University of Alberta

Perception and Communication – Their Relationship in Collective Sorting

by

Sean R. Verret

A thesis submitted to the Faculty of Graduate Studies and Research in partial fulfillment of the requirements for the degree of **Master of Science**.

Department of Electrical and Computer Engineering

Edmonton, Alberta
Fall 2004

Library and Archives Canada

Published Heritage Branch

395 Wellington Street
Ottawa ON K1A 0N4
Canada

Bibliothèque et Archives Canada

Direction du Patrimoine de l'édition

395, rue Wellington
Ottawa ON K1A 0N4
Canada

Your file Votre référence
ISBN: 0-612-95872-8
Our file Notre référence
ISBN: 0-612-95872-8

The author has granted a non-exclusive license allowing the Library and Archives Canada to reproduce, loan, distribute or sell copies of this thesis in microform, paper or electronic formats.

L'auteur a accordé une licence non exclusive permettant à la Bibliothèque et Archives Canada de reproduire, prêter, distribuer ou vendre des copies de cette thèse sous la forme de microfiche/film, de reproduction sur papier ou sur format électronique.

The author retains ownership of the copyright in this thesis. Neither the thesis nor substantial extracts from it may be printed or otherwise reproduced without the author's permission.

L'auteur conserve la propriété du droit d'auteur qui protège cette thèse. Ni la thèse ni des extraits substantiels de celle-ci ne doivent être imprimés ou aturement reproduits sans son autorisation.

In compliance with the Canadian Privacy Act some supporting forms may have been removed from this thesis.

Conformément à la loi canadienne sur la protection de la vie privée, quelques formulaires secondaires ont été enlevés de cette thèse.

While these forms may be included in the document page count, their removal does not represent any loss of content from the thesis.

Bien que ces formulaires aient inclus dans la pagination, il n'y aura aucun contenu manquant.

Canada

Science is nothing but perception.
Plato (427 BC - 347 BC)

To my parents and family
for all their encouragement.
Without your support none of this could
have ever been accomplished!

Acknowledgements

I would like to take this opportunity to express my gratitude to those that have supported my research, because without their support this work would not have been possible. I have been very lucky to have Dr. Max Meng and Dr. Hong Zhang as my co-supervisors. Both have been great mentors and great friends. Hong has especially encouraged and supported me during times of adversity and promoted me to explore realms of research that I wouldn't have been able to achieve. Additionally, I am very grateful to my committee members, Dr. Mrinal Mandal and Dr. Michael Bowling who took the time to read and comment on my work. Finally, I am indebted to NSERC and AMC Technologies for providing me with the funds necessary to carry out my research and studies.

I would like to give a special thanks to Joan Digney for assisting me with the technical aspects of writing my thesis and for helping me with all of my grammerical mistakes. Thank you for magnifying my perceived intelligence!

A plethora of BEvERages go out to my friends Matthew McNaughton and Andrzej Zadorozny. Without their tireless efforts in the lab working on RoboCup and the entire system, none of this work would have been possible. They were always available whenever I needed help and I'm grateful for it. Additional BEvERages go out to the rest of the RoboCup members from the past 4 years especially Jon Kelly, Doug Kondor, Chris Parker, Axel von Bertoldi, and the ECE Machine Shop.

One last thanks goes out to Dr. Simon Monckton and Dr. Bruce Digney who both provided additional guidance when I was off campus and to Defence Research and Development Canada for giving me the time to complete my writing.

Finally, I'd like to thank my parents, family and friends for their unwavering encouragement, nagging, pestering, questioning, harassment, badgering, etc. Without that "support" I would not have been able to finish in the short time (4+ years) it took me.

Table of Contents

1 Introduction 1

2 State of the Art in Collective Robotics 3
- 2.1 Introduction . . . 3
- 2.2 Before Collective Robotics . . . 3
  - 2.2.1 Traditional Robotics . . . 4
  - 2.2.2 Behaviour Based Robotics . . . 5
  - 2.2.3 Why Behaviour Based Robots . . . 6
- 2.3 Collective Robotics . . . 7
  - 2.3.1 What is Collective Robotics . . . 7
  - 2.3.2 Why Collective Robotics . . . 7
  - 2.3.3 Collective Robotics Examples . . . 8
  - 2.3.4 Biological Influence . . . 11
- 2.4 Summary . . . 12

3 Research Problem 13
- 3.1 Introduction . . . 13
- 3.2 Task Selection . . . 13
  - 3.2.1 Work Done in Multi-Robot Sorting . . . 14
  - 3.2.2 Work Done in Multi-Robot Communications . . . 15
  - 3.2.3 What Has Not Been Done? . . . 16
- 3.3 Why? . . . 16
- 3.4 Thesis Statement . . . 16
- 3.5 Summary . . . 16

4 Theory 18
- 4.1 Introduction . . . 18
- 4.2 Perception . . . 18
  - 4.2.1 Necessity/Relevance . . . 19
  - 4.2.2 Definition . . . 19
- 4.3 Communication . . . 20
  - 4.3.1 Necessity/Relevance . . . 20
  - 4.3.2 Definition . . . 21
- 4.4 Sorting Method . . . 22
  - 4.4.1 Necessity/Relevance . . . 22
  - 4.4.2 Definition . . . 23
  - 4.4.3 Maximum Density . . . 25
- 4.5 Efficiency Calculation . . . 26
  - 4.5.1 Definition . . . 26
- 4.6 Summary . . . 27

5 Experimental Setup **28**
5.1 Introduction . . . 28
5.2 Assumptions . . . 28
5.3 Robots . . . 29
5.3.1 Mechanical . . . 29
5.3.2 Electrical . . . 29
5.3.3 Firmware . . . 33
5.3.4 Complete Robot . . . 34
5.4 Environment . . . 34
5.4.1 Vision System . . . 34
5.4.2 Wireless Communications . . . 35
5.4.3 Robot Monitoring . . . 35
5.5 Methods . . . 35
5.6 Summary . . . 37

6 Experimental Results **38**
6.1 Introduction . . . 38
6.2 Single-Robot Experiments . . . 38
6.3 Single-Robot Experiment Data . . . 39
6.4 Multi-Robot Experiments . . . 40
6.5 Multi-Robot Experimental Data . . . 41
6.6 Summary . . . 48

7 Discussion **49**
7.1 Introduction . . . 49
7.2 Effect of Initial Conditions . . . 49
7.3 Effect of Population Size . . . 52
7.4 Effect of Environment . . . 52
7.5 Effect of Sensing and Communication Range . . . 53
7.5.1 Sensing Range . . . 53
7.5.2 Communication Range . . . 53
7.6 Experiment Specific Tendencies . . . 54
7.7 The Stopping Criterion . . . 55
7.8 Similarities to Other Multi-Robot Systems . . . 56
7.9 Summary . . . 56

8 Conclusions **57**
8.1 What has been done? . . . 57
8.1.1 Successful Segregation Sorting Algorithm . . . 57
8.1.2 Communication Improves Task Performance . . . 57
8.1.3 Communication is Related to Perception . . . 58
8.2 Importance . . . 58
8.3 Future Research . . . 59

Bibliography **60**

List of Tables

6.1 Time (in minutes) taken by 1, 2 or 4 robots, with and without communication present, to sort objects of two classes into two piles with sensing ranges of 50 cm, 100 cm and 180 cm. 39

List of Figures

2.1 A traditional or hierarchal decomposition of a mobile robot control system into functional modules [12]. 4
2.2 A decomposition of a mobile robot control that is based on task achieving behaviours [12]. 5
2.3 Clustered objects. 9
2.4 Segregated objects. 9
2.5 Patch sorted objects. 9
2.6 Annular sorted objects. 9

4.1 A sample world, W. The white and black circles are objects of type 1 and 2. The larger white circles with black circles inside are robots and the arrows in the circle signify their orientation. The triangular areas coming from a robot is its area that is sensed s_r. 19
4.2 Communicating robots (R1 can perceive what R2 can perceive). 21
4.3 Sorting algorithm state machine 23
4.4 Honeycomb pattern 25
4.5 This figure illustrates Equation 4.14. Given one pile with three objects, one pile with five objects and the last pile with two objects, the percent complete is $(\frac{3}{10})^2 + (\frac{5}{10})^2 + (\frac{2}{10})^2 = 38\%$ complete. 27

5.1 The main controller board used in our robots. The controller is based around the MC68332. It contains 512K of Flash, 512K of RAM and measures 2.95 inches square. 30
5.2 An example of how various PWM duty cycles are produced by XORing two phase-shifted pulse trains using the MC68332 MCPWM routines. 30
5.3 The wireless communications board we used in our robots. The RPC-418-40 is designed and manufactured by Radiometrix. 31
5.4 The motor board used in our robots. This board contains all the high current components and controller components necessary to drive a robot's motors. . 32
5.5 The IR break beam board used in our robots. This board simply two voltages and returns a 1 if the compared voltage is higher than the reference voltage or a 0 if the compared voltage is lower than the reference voltage. 32
5.6 A robot motor. This motor is able to deliver full quadrature decoding signals. 33
5.7 The PD controller used to control the motor velocities. The K_p and K_d constants were tuned experimentally. V_a is the actual velocity as sensed by the motor encoder and V_d is the desired velocity specified to the controller. . 33
5.8 One of the robots used in our experiment along with an object that is being controlled by the robot. 34
5.9 This is the base station used to transmit signals to and receive signals from the robots. 35
5.10 A screen shot of BotCommander. BotCommander (the robot monitor) was used to monitor the robots and objects in the experiments. The experiments could be started or stopped with a click of the mouse. 36

5.11 Initial arrangement of the objects (24 objects for each class), which are uniformly distributed and perfectly mixed. Note: There is considerable radial distortion in the camera causing the rectangular field to appear to be barrel-shaped. Distortion correction was performed before the coordinates of the objects were calculated. 37

6.1 These graphs represent the change in percent complete versus time for a single-robot experiment in which the robots has a sensing range of 50 cm. . . 39
6.2 These graphs represent the change in percent complete versus time for a single-robot experiment in which the robot has a sensing range of 100 cm. . . 40
6.3 These graphs represent the change in percent complete versus time for a single-robot experiment in which the robot has a sensing range of 180 cm. . . 40
6.4 Snapshots of the experiment involving four robots, each with a 50 cm sensing and communication range. Note: There are two types of objects (yellow and orange) as in Figure 5.11, except that yellow objects have been painted white with an image editor to improve the readability of the images. (a) shows the initial conditions (b) shows the experiment status after 8 minutes, (c) shows the experiment status after 15 minutes (d) shows the experiment status after 23 minutes (e) shows the experiment status after 30 minutes (f) shows the final results of the experiment. 41
6.5 These graphs represent the change in percent complete versus time for a two-robot experiment in which the robots have a sensing range of 50 cm and are not able to communicate with each other. 42
6.6 These graphs represent the change in percent complete versus time for a two-robot experiment in which the robots have a sensing range of 100 cm and are not able to communicate with each other. 42
6.7 These graphs represent the change in percent complete versus time for a two-robot experiment in which the robots have a sensing range of 180 cm and are not able to communicate with each other. 43
6.8 These graphs represent the change in percent complete versus time for a two-robot experiment in which the robots have a sensing range of 50 cm and are able to communicate with each other. 43
6.9 These graphs represent the change in percent complete versus time for a two-robot experiment in which the robots have a sensing range of 100 cm and are able to communicate with each other. 44
6.10 These graphs represent the change in percent complete versus time for a two-robot experiment in which the robots have a sensing range of 180 cm and are able to communicate with each other. 44
6.11 These graphs represent the change in percent complete versus time for a four-robot experiment in which the robots have a sensing range of 50 cm and are not able to communicate with each other. 45
6.12 These graphs represent the change in percent complete versus time for a four-robot experiment in which the robots have a sensing range of 100 cm and are not able to communicate with each other. 45
6.13 These graphs represent the change in percent complete versus time for a four-robot experiment in which the robots have a sensing range of 180 cm and are not able to communicate with each other. 46
6.14 These graphs represent the change in percent complete versus time for a four-robot experiment in which the robots have a sensing range of 50 cm and are able to communicate with each other. 46
6.15 These graphs represent the change in percent complete versus time for a four-robot experiment in which the robots have a sensing range of 100 cm and are able to communicate with each other. 47

6.16 These graphs represent the change in percent complete versus time for a four-robot experiment in which the robots have a sensing range of 180 cm and are able to communicate with each other. 47

7.1 Four examples of sorting experiments shortly after the experiment has been started. (a) 2 robots with a sensing range of 180cm and no communication after 8 minutes. (b) 4 robots with a sensing range of 100cm and no communication after 8 minutes. (c) 4 robots with a sensing range of 100cm and with communication after 9 minutes. (d) 4 robots with a sensing range of 50cm and with communication after 10 minutes. Note: There are two types of objects (yellow and orange) as in Figure 5.11, except that yellow objects have been painted white with an image editor to improve the readability of the images. 50

7.2 Four examples of sorting experiments after the experiment have completed. (a) 2 robots with a sensing range of 180cm and no communication. (b) 4 robots with a sensing range of 100cm and no communication. (c) 4 robots with a sensing range of 100cm and with communication. (d) 4 robots with a sensing range of 50cm and with communication. Note: There are two types of objects (yellow and orange) as in Figure 5.11, except that yellow objects have been painted white with an image editor to improve the readability of the images. 51

7.3 Completion times for each of the three robot numbers with 180cm sensing range when there is either communication (bottom curve) or no communication (top curve) between the robots. 52

7.4 Completion times for each of the three sensing ranges with two robots when there is either communication (bottom curve) or no communication (top curve) between the robots 54

7.5 Completion times for each of the three sensing ranges with four robots when there is either communication (bottom curve) or no communication (top curve) between the robots 54

7.6 This graph represents the change in percent complete versus time for a four-robot experiment in which the robots have a sensing range of 100 cm and are able to communicate with each other. Notice the sharp increase in the slope of the curve at the end of the graph. Several experiments showed results like this because there were fewer piles to drop objects onto. 55

7.7 This graph represents the change in percent complete versus time for a four-robot experiment in which the robots have a sensing range of 100 cm and are not able to communicate with each other. Notice the dip in performance around the 33 minute mark. This indicates a time when a robot made an incorrect decision and accidentally separated a pile of objects into more than one pile. 55

8.1 As can been seen in this image, one of the objects of type 1 has been stranded in the middle of the type 2 objects. The completion criteria was still satisfied, however, the experiment didn't finish with a 100% completion rate. This experiment involved 2 robots with a sensing range of 180 cm and no communication between the robots. 58

List of Symbols

W	world/environment
s_i	portion of W sensed by robot i
${}^{W}v$	a perceived object in W
t	type of object (1= type 1 or 2= type 2)
p^t	object of type t
C_i	communicated data to robot i
$\mathcal{P}_i$	perceived data of robot i
$\mathcal{P}'_i$	perceived data of robot with communication i
n_t	number of centroids of type t
n_r	number of robots
N_t	number of visible objects of type t
N_r	number of visible robots t
m	the number of objects viewable
c_g	center coordinate that is central to all viewable objects
d	distance from one object to c_g
d_T	total distance of all objects to c_g
ρ	density of objects in a region
ρ^*	the density of the most compact cluster of m objects
x	x-coordinate of an object in W
y	y-coordinate of an object in W
r	robot
θ	orientation of object
κ	total number of visible objects
δ	robot's sensing distance
ω	robot's sensing angle
$\mathcal{U}$	pick up value between 0 and 1
$\mathcal{D}$	drop value between 0 and 1
A_i	the number of objects in cluster i
A_t	the number of objects in W
α	constant
K_p	proportional constant
K_i	integral constant
K_d	differential constant

Chapter 1

Introduction

Do you have glasses? Are you near-sighted? Have you ever participated in a scavenger hunt? Were you a member of a team or just an individual? Imagine that you're very near-sighted and don't have your glasses on. Can you imagine participating in a scavenger hunt where you had to find one hundred items all by yourself? Now put on your glasses. Will it take less time to find the items? Now imagine that you have three or four friends on your team. Do you think you can find all the items in a faster, more efficient manner? Now imagine that you and your friends all have cell phones. Can you find all the items even faster?

Competing in a scavenger hunt without your glasses is indeed possible as an individual. Putting on your glasses only improves your chances. When you join with a team it's most likely going to improve the time it takes to find all the items. Finally, if everyone on your team can communicate to each other about what they've found and where certain items can be found, it should take even less time to find everything.

Can we apply these same basic principles to robotics? Can we make robots "see" better to improve their efficiency? Can we simply add more robots to a task to improve the overall efficiency? Can we enable communication among these robots to further improve their efficiency? How do all of these variables relate to each other?

In this thesis, we will describe a very similar scenario. We will describe a collective sorting system. In other words, we will be talking about a team of robots that have been designed to coordinate and cooperate with each other to sort groups of objects into separate piles of objects of the same type. The collective sorting problem is a canonical problem in collective robotics [14]. Past efforts have employed a minimalist approach with minimum sensing, communication, and memory. In addition, segregation sorting, in which the number of clusters is exactly the number of classes N, has yet to be achieved experimentally.

In this system we will be able to vary the sensing range of the individual robots, and control whether or not the robots can communicate with each other. We will also vary the number of robots participating in the experiments.

In this thesis, we describe a sorting algorithm that has successfully completed segre-

gation sorting *experimentally* nearly 100% of the time, under the assumption of a sensing range greater than that in previous studies. In addition, we examine the relationship between sensing and communication, and how the task performance changes with sensing and communication ranges, for different sizes of robot populations. As in previous studies, our algorithm remains simple rule based, where decisions on whether an object should be moved are determined with only locally sensed or communicated information[1].

Before tackling the research problem we'll first discuss the state of the art (Chapter 2) and specifically define the research problem (Chapter 3). In the next chapters we'll discuss the theory (Chapter 4) behind the experiments (Chapter 5) and the results (Chapter 6). Finally we'll finish off with a discussion (Chapter 7) and our conclusions (Chapter 8).

[1] Portions of this thesis have been published. Sean R. Verret, Hong Zhang and Max Q.-H. Meng. Collective Sorting with Local Communication. In *IEEE Conference on Intelligent Robots and Systems IROS-04*, 2004

Chapter 2

State of the Art in Collective Robotics

2.1 Introduction

The word "robot" was first introduced by the acclaimed Czech playwright Karel Capek (1890-1938). Derived from the Czech word "robota" for forced labor or serf, the word "robot" was introduced into his play R.U.R. (Rossum's Universal Robots) which opened in Prague in January 25, 1921 [50]. Since then, "robot" has been defined in several different contexts and the definition will undoubtedly continue to evolve in the future. The purpose of this chapter is to give the reader some brief background on collective robotics. Several examples of collective robotics are given and emphasis is put on experiments specifically dealing with collective sorting and the communication challenges involved in collective robotics. First, a very brief summary of the robotics world before collective robotics was explored will be given.

2.2 Before Collective Robotics

In robotics several terms are used in different contexts. For this thesis several of these terms will be used as described below.

"Robot," according to Webster [47], will be defined as an automatic device that performs functions normally ascribed to humans or a machine in the form of a human.

"Mobile Robot," according to a collection of sources, will be defined as a robot capable of changing its base coordinate frame with respect to a global, stationary coordinate system. More specifically, a robot with nonholonomic constraints to its environment. Examples include wheeled, walking, hopping, squirming, swimming, and flying robots.

"Autonomous Robot" will be defined as a robot that is not controlled by any human sources or by human forces; i.e. it is a robot that acts independently of human control.

"Behaviours" are a direct mapping of sensory inputs to a pattern of motor actions that

are then used to achieve a task [50], or more simply, behaviours are the transfer function between sensing and acting.

2.2.1 Traditional Robotics

The first actual use of a robot was in 1942, when the United States government undertook a project, called the Manhattan Project, to build a nuclear bomb [50]. The telemanipulator was first developed when the scientists involved came across problems with the handling and processing of radioactive materials. The telemanipulator was handled by the scientist who would move it around and watch a display of what the end of the manipulator was doing. This is very similar to what is now being done in the field of medicine as robotic telemanipulators are helping surgeons perform key maneuvers without taking drastic measures on the patient. It was the partial automation of this first manipulator that paved the way towards automated manipulators for other applications.

It is widely agreed that the first programmable robot was designed by George Devol, who developed Unimation. Unimation [50] was bought by General Motors in the early 1960's and was the first robot of its kind to be used in industry. Since then the two most common types of robot technology that have evolved for industrial use are robot arms (industrial manipulators) and mobile carts (automated guided vehicles) [50]. However, current research is developing teleoperated robots equipped with sensors that acquires information and communication links that sends sensory data to the operator. Eventually, robots that acted semi-autonomously or under supervisory control were being developed.

In 1967, Shakey, the first Artificially Intelligent (AI) robot was created at the Stanford Research Institute. Shakey operated using a strictly hierarchical pattern of intelligence as shown in Figure 2.1. A hierarchical pattern of intelligence operates sequentially and in an ordered fashion. A step cannot be initiated until the step before has finished. Until the late 80's, the hierarchical approach, was "the" method used by researchers.

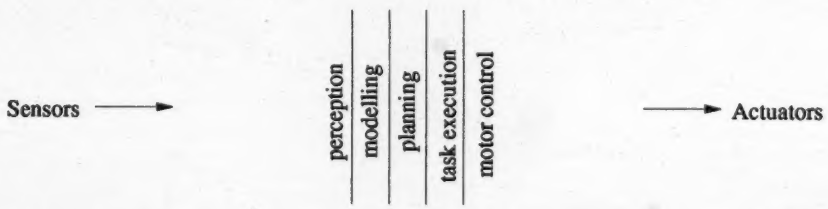

Figure 2.1: A traditional or hierarchal decomposition of a mobile robot control system into functional modules [12].

It is important to understand, however, that even today, the majority of robots reside in highly structured worlds and very closely monitored environments that are amenable to traditional AI approaches.

2.2.2 Behaviour Based Robotics

Valentino Braitenberg wrote a book [11] that described a set of experiments in which increasingly complex vehicles are built from simple mechanical and electrical components. Each of these imaginary vehicles in some way mimics intelligent behaviour and each was labeled with a name corresponding to the behaviour it imitates, i.e., Fear, Aggression, Logic, Values, etc. Braitenberg progresses and the reader sees very intricate behaviours emerging from the interaction of simple component parts. Essentially, by the end of the book Braitenberg builds intelligent behaviour.

Until 1986, intelligence was implemented using the hierarchical approach mentioned in Section 2.2.1. At about the same time as Braitenberg was developing his hypothetical experiments, Rodney A. Brooks wrote his paper titled, "A Robust Layered Control System for a Mobile Robot" [12]. The entire mobile robotics scene was revolutionized. Brooks outlined a layered, behaviour based approach (see Figure 2.2) that acted in a nearly orthogonal fashion to the hierarchical approach and very similarly to the approach described by Braitenberg's experiments [11]. Brooks outlined a system that was designed solely for the purpose of mobile robotics. He outlined the requirements of a control system for an intelligent autonomous mobile robot that could possibly include: multiple goals, multiple sensors, robustness, and additivity [12]. These requirements put constraints on any possible future control system. Brooks came up with a method, largely labeled the "behaviour based approach", that simplified the control system needed for autonomous mobile robots. This system paved the way for mobile robots that could be created, maintained and improved upon in a more efficient manner. The reason being that if a new sensor was added or needed repairs, or if more capabilities were needed, a new behaviour could be added or dropped with ease.

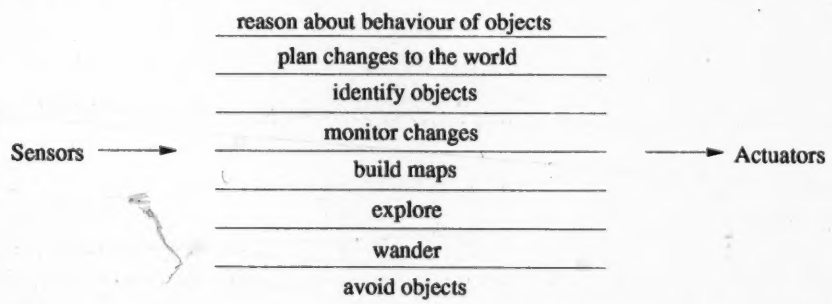

Figure 2.2: A decomposition of a mobile robot control that is based on task achieving behaviours [12].

Around the same time, Reynolds [57] was examining flocks of birds and the aggregate motion associated with a flying flock. Reynolds developed a distributed behavioural model

for a bird in a flock. His approach assumed that a flock is simply the result of the interaction between the behaviours of the individual birds. Reynolds created simple behaviours like collision avoidance, velocity matching, and flock centering into his simulated birds. The approach developed by Reynolds was very similar to Brooks' methods mentioned above.

These three very separate sources eventually formed what is now known as behaviour based robotics. Brooks has also provided additional work [13] into behaviour based systems as have others [27, 38].

2.2.3 Why Behaviour Based Robots

Why did behaviour based robots appear? Why are behaviour based robots "simpler" and what are the consequences of these simpler robots?

There are several reasons why behaviour based robotics has emerged in the field of mobile robots.

- Robots need to execute quickly and a tighter coupling with the sensors and the actuators allow the robots to "react" and operate in near "real-time." The overall simplicity of the behaviour based system means that such systems have excellent real-time performance, even with modest resources [10].
- Behaviours can be implemented in either software or hardware with very low complexity and thus can be executed quickly regardless of computational power.
- Behaviours are controlled by the environment; in the fact that they operate reactively, rather than by storing previous states in memory and then acting accordingly.
- Behaviour based robots can be layered with incremental bits of circuitry. Circuitry is used here loosely, meaning either direct physical circuitry or firmware organized in a manner similar to circuitry. By leaving old circuitry (layers) in place, there is an ability to continue to operate even if the new circuitry fails [13].
- Behaviours are inherently modular and easy to test in isolation from the system [50].

However with these robots there are some consequences:

- Robots aren't necessarily robust. I.e., if an upper layer fails there are usually no mechanisms to indicate that a degradation has occurred [50].
- Finite state mechanisms are only as good as the level of situations that can be anticipated or incorporated into the state machine [50]. I.e., hardwired behaviours result in the robot not being able to adapt to new unforeseen situations
- As with humans, reactive robots will react according to the environment, but won't necessarily do the correct thing [50].

- Most behaviour based robots lack a planning/reasoning component – they cannot predict consequences of actions. However, there is work being done so that new behaviours can be learned using learning techniques such as neural networks, Bayesian nets and reinforcement learning. Also, researchers are inserting global knowledge and planning abilities on a higher level deliberative system, which lies on top of behaviour-based system

As mobile robots became easier to create, researchers then started incorporating the research that was going on in the biological sciences into controlling groups of robots and collective robotics was born.

2.3 Collective Robotics

2.3.1 What is Collective Robotics

Collective robotics is quickly becoming a vast research area and includes several different topics and ideas, as shown in the various surveys [2, 17, 14, 54].

Before we begin analyzing these fields, we need to define what collective robotics is and define a few other terms that will be used often in this thesis. Cao *et al.* [14] state that cooperative behaviour "is a subclass of collective behaviour that is characterised by cooperation." In other robotics literature, cooperation has been defined in several different ways including:

1. "Joint collaborative behaviour that is directed toward some goal in which there is a common interest or reward" [9].
2. "A form of interaction usually based on communication" [42].
3. "[Joining] together for doing something that creates a progressive result such as increasing performance or saving time" [55].

After we've examined all the different definitions of cooperation we need to define cooperative behaviour. Cao *et al.* [14] define cooperative behaviour as a system that, due to some underlying mechanism of cooperation, increases the total utility of the system.

Now that we've come up with an idea of what collective robotics is, we need to understand why collective robotics is an important area to study.

2.3.2 Why Collective Robotics

Dudek *et al.*, explained that collectives offer the possibility of enhanced task performance, increased task reliability and decreased cost over more traditional robotic systems [17].

Cao *et al.* [14] says that systems of multiple robots are interesting because the overall task may be too complex, and/or building and using several simple robots can be easier,

cheaper, more flexible and more fault tolerant. Cao *et al.* [14] also discusses how a single robot is spatially limited when "multiple robot systems can accomplish tasks that no single robot can accomplish, since ultimately a single robot, no matter how capable, is spatially limited."

Arkin and Balch [3] argued that two (or more) robots can be better than one for several reasons:

- Many robots can be in many places at the same time (distributed action).
- Many robots can do many, perhaps different things at the same time (inherent parallelism).
- Often each agent in a team of robots can be simpler than a more comprehensive single robot solution (simpler is better).

Cao *et al.* [14] discuss the wide range of disciplines involved with collective robotics such as artificial intelligence, game theory/economics, theoretical biology, distributed computing/control, animal ethology and artificial life.

2.3.3 Collective Robotics Examples

We stated earlier (Section 2.2.1) that the birth of robotics occurred in 1942. However, as early as the 1940's, collective robotics was being born with the research of Grey Walter and his turtle-like robots equipped with light and touch sensors. These robots were some of the first to interact with each other by responding to each other's movements [63, 64]. Elsie and Elmer were not widely publicized, perhaps because Walter didn't call them robots and though they may not have provided the inspiration for collective robotics, they are indeed the first example of robots that could coordinate and coexist. Since Elsie and Elmer, there have been several different task domains that researchers have used to demonstrate collective robotics. This list is by no means exhaustive, but does include a wide range of different environments.

- **Foraging/Sorting:** Clustering of objects into piles by homogeneous and heterogeneous groups of robots has been examined in many different ways by several different researchers. Initially, simply clustering objects into a pile, as demonstrated by [10, 39, 40], was a problem that was overcome. Similarly, the trash-collecting robots [8] of Georgia Tech and the artificial toxic waste cleanup mission [53], described by Parker, both looked into the cleaning up of specific areas.

  Much of this research was most likely inspired by the writing of Deneubourg *et al.* [16] in their paper "The Dynamics of Collective Sorting Robot-like Ants and Ant-like

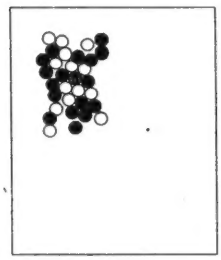

Figure 2.3: Clustered objects.

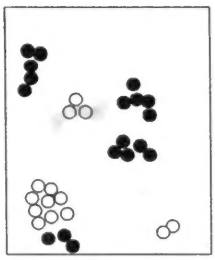

Figure 2.4: Segregated objects.

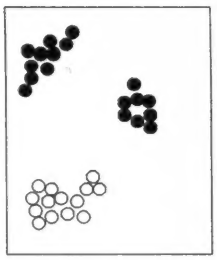

Figure 2.5: Patch sorted objects.

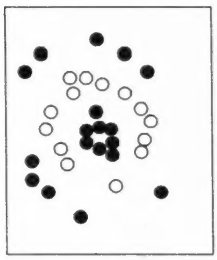

Figure 2.6: Annular sorted objects.

Robots." Deneubourg *et al.* describe a simple behaviour algorithm that is used by each worker and generates a sorting process.

Eventually, Melhuish *et al.* [46] built a physical system using minimalist agents that were able to sort colored Frisbees into annular rings. Melhuish *et al.* also went on to differentiate four different types of sorting [46]:

1. Clustering: grouping a class of objects within a contiguous area that is a small fraction of the area of the available environment. See Figure 2.3.
2. Segregation: grouping two or more classes of objects so that each occupies a continuous area of the environment that is not occupied by members of any other class. See Figure 2.4.
3. Patch sorting: grouping two or more classes of objects so that each is both clustered and segregated, and each lies outside the boundary of each other. See Figure 2.5.
4. Annular sorting: forming a cluster of one class of objects, and surrounding it with annular bands of the other classes, each band containing objects of only one type. See Figure 2.6.

Finally, Lerman [37] studied the effects of interference in a group of foraging robots and examined the overall effectiveness of a group of robots as their numbers increased in a similar-sized environment.

- **Multi-Robot Communication:** The problem of direct inter-robot communication is an important subject of research in collective robotics. There is a general consensus among the researchers who have examined this problem with regard to the role of communication in MRS cooperative tasks. Balch and Arkin [6] ran experiments with LED indicators on top of the robots signifying what state they were in. They concluded that communication significantly improves performance in tasks with little environmental communication, and that complex communication strategies offer little or no

benefit over low-level communication. Dahl *et al.* [15] ran experiments with robots attached with GPS modules that enabled communication between robots that were within $2m$ of each other and observed that individual controllers using communication are significantly more effective than those using a communication-free controller. Easton and Martinoli [18], using infrared signaling on their Khepera robots, found that communication schemes heavily influence the rate of success of the system. Finally, Matarić [43] used communication to share sensory data to overcome a hidden state and share reinforcement to overcome the credit assignment problem between the agents. She concluded that communication may compensate for the limitations of more direct sensory modalities, thus enabling efficient learning in multi-robot systems.

Researchers in this field generally agree that adding communication to the system leads to an improvement in performance, with respect to many different task-dependent performance metrics.

- **Cooperative Transport:** The cooperative transport of large objects (especially objects that are much bigger than the robots) is particularly interesting because a single robot is unable to accomplish this task. Kube's experiments [36, 35] provided the first formalized model of cooperative transport in ants with his wheeled robots. Matarić [44] went on to do similar work with her 6-legged box-pushing robots. The cooperative transport problem was again examined by Trebi-Ollennu *et al.* [62] when they discussed a Mars Rover pair cooperatively transporting a long payload. Other research in cooperative transport includes [48].

- **Landmine Detection:** Landmine detection has been studied for many years. Many times the problem has been solved by very sophisticated machines with a wide array of sensors available to the machine. Franklin *et al.* [19] were some of the first researchers to tackle this problem with distributed sensing and multiple search agents. Gage's [23] research outlined how the evolution of today's sensors will soon make it possible for inexpensive yet highly effective autonomous agents to accomplish such tasks as landmine detection.

- **Collective Building:** Collective building or collective construction is a rather new area of exploration in the collective robotics realm. Most research that has been done has been biologically inspired by ant and termite collectives. These researchers have also attempted to complete their construction tasks using robots with minimalist behaviours. Wawerla *et al.* [65] may have been the first to demonstrate collective construction abilities in robots. The also demonstrated the ability of robots to construct simple 2D structures. In similar work, Parker [52, 51] designed a group of minimalist robots that could perform blind bulldozing techniques for the clearance of an area

mimicking a specific species of ant, the Leptothorax tubero-interruptus.

- **Traffic Control:** Cao *et al.* [14] describe the traffic-control task domain or collision avoidance domain as a problem of resource conflict, which may be resolved by introducing traffic rules. These traffic rules are commonly used to reduce planning cost for avoiding collisions and deadlock in real-world environments, such as a network of roads. Nearly all multi-robot systems have to deal with obstacle avoidance, path planning etc., although Cao *et al.* [14] state that in the experiments that use the behaviour-based approach, robots are never restricted to road networks.

- **Military/Space Applications:** Huntsberger and his colleagues have written several different papers on the robotic challenges of developing multiple planetary rovers for different types of missions [33, 34, 32]. Balch [7] has also written about different military based applications in which behaviour-based formation control of multi-robot teams would be useful.

- **Robot Formation:** As mentioned above Balch [7] has studied military-based formation strategies. Reynolds [57], however, may have been one of the first to study flocks and their formations with his distributed behaviour model back in 1987. Since then Tang and Zhang [61], as well as Fredslund and Matarić [22], have also performed separate studies of robot formation in a decentralized environment.

- **Collective Mining:** Dunbar and Klein [60] wrote a technical note in 2002 that considered using mini- or micro-machines to improve control and efficiency of fragmentation, heap leaching and other mining and mineral processing operations.

- **Collective Exploration:** Arkin and Diaz [4] discussed a line-of-sight constrained exploration algorithm for reactive multi-agent robotic teams.

2.3.4 Biological Influence

Many AI roboticists often turn to the biological sciences for inspiration because animals and man provide existence proofs of different aspects of intelligence [50]. The proof that intelligence exists has driven these roboticists to attempt to mimic this intelligence in machines. Insects are considered a very "simple" organism but yet they are still able to accomplish intelligent tasks despite the fact that they have basically no sophisticated intelligence. Thus, nature has provided the robotics community with an abundance of different "behaviours" that designers have attempted to incorporate into their designs.

The biological influence that definitely appears in collective robotics is not seen as a proof of collective robotics but rather as an inspiration for collective robotics. Some of the

most important biological species that have influenced the collective robotics research are those of insects such as ants [48, 1, 24], bees, wasps, termites [49], etc.

One of the more comprehensive references on insects is the book by Holldobler and Wilson [31], which gives an excellent reference on the communication and social organization between ants. They also discuss how the communication between ants is typically of the chemical (pheromones) variety. Others, like Deneubourg [16] and Franks [21, 20] have written publications about foraging and sorting amongst ants.

There are also other notable works that have discussed biological inspired experiments in robotics [41, 24, 1].

These researchers and others have laid out simple behaviour-based models that can accomplish complex tasks such as foraging and cooperative transport. Researchers [36, 35, 10, 46] have then gone on to mimic these behaviours with real robots and have been able to prove experimentally that the use of basic behaviours inspired by the biological world can reproduce similar behaviours in the robotic world.

The biological world has indeed inspired the collective robotics community to ask important questions about how emergent behaviours and emergent intelligence actually occur, especially amongst "simple" organisms.

2.4 Summary

As we've seen in this section, there are several areas of research currently being explored in the field of collective robotics. In the next chapter, we will discuss more specific results from researchers in the field of collective sorting and multi-robot communications. This will lead us to the fundamental research problem being examined by this thesis.

Chapter 3

Research Problem

3.1 Introduction

As was mentioned in the Chapter 2, there have been a vast assortment of tasks and experiments done in the area of collective robotics research. This research has included but are not limited to: traffic, cooperative transport, foraging, landmine detection, collective building, robot formation, etc.

This chapter will explain why we chose the sorting task as our problem environment and we'll then go into more detail than Section 2.3.3 about the collective sorting problem and how it has been tackled. In the next section we will discuss the multi-robot communication problem in more detail. Finally, we'll explain why this research is interesting and conclude with our thesis statement.

3.2 Task Selection

As explained in Section 2.3.3, the sorting problem has been examined and researched. This has been an interesting problem in the collective robotics area since it can be solved with one robot. However, by introducing more robots in to the environment, researchers have examined how the process of sorting can be made more efficient. Efficiency, of course, is defined in this instance as an increase in speed at which the desired goal of a sorted environment may be accomplished.

It has been mentioned by Cao *et al.* [14] that the sorting/foraging task is interesting because it can be performed by each robot independently. However, the issue in the collective robotics field is whether adding multiple robots reflects a performance gain on the entire system. Performance gain or efficiency can be measured in several different ways and our method of efficiency will be discussed in Section 4.5. Cao *et al.* [14] also noted that a the foraging/sorting task becomes trivial if communication is involved. Thus, based on the triviality of the task of sorting we decided to use the sorting environment as a platform for our research.

3.2.1 Work Done in Multi-Robot Sorting

Multi-robot sorting has been physically performed by several different researchers. Most notably, the work done by Beckers *et al.* [10] and Melhuish *et al.* [46] are two concrete examples of robots sorting different objects into respective piles.

Beckers *et al.* [10] designed robots that would gather 81 randomly placed objects and cluster them into a pile. Beckers' clustering methodology with robots is very similar to the observations of Deneubourg *et al.* [16], where similar clusters are observed in ant colonies, generated by the probabilistic behaviour of the workers. Although, no direct evidence that ants employ probabilistic rules has been discovered. Beckers' *et al.* created behaviour based robots that could cope well with an unstructured environment and were inexpensive. It was their assumption that the biological principal of stigmergy would fit better with the biologically inspired architectures of behaviour based robots than with the "alien" computational paradigm of conventional robotics [10].

Beckers *et al.* implemented a three behaviour system. These three behaviours were move straight, avoid obstacles and drop object. The robot would be in the move straight behaviour for most of the time. Once it had gathered three objects, it would then shift into the drop behaviour. Also at any time that the robot's touch sensors were engaged was the avoid behaviour triggered. The robots, within this world, operated autonomously and independently. All of the robots' sensors, motors and control hardware were on board and there was no explicit intra-robot communications [10]. From the literature reviewed, it appears that at the time of Beckers' research there hadn't been any robotic examples of clustering.

Melhuish *et al.* [46, 30] designed robots that were able to mimic brood sorting in ants. Melhuish *et al.* also described the different types of sorting as was mentioned in Section 2.3.3. Melhuish's sorting method included no capacity for spatial orientation or memory. As with many researchers (Section 2.3.4), Melhuish's methods were inspired by the biological world. Melhuish quotes both Deneubourg's and Franks' works with different species of ants including the *Leptothorax unifasciatus*.

Melhuish *et al.* [46] examined several different methods of clustering and sorting. They were able to demonstrate clustering using similar rules as [10], but went on to further the research by adding differences to the "drop" behaviour. Melhuish was able to create annular sorted rings of different colors of objects simply by varying the drop distance that the robot would pull back based on the color of the object it was holding.

Based on the literature that has been reviewed to date it appears that at the time of Melhuish's research there hadn't been any robotic examples of sorting similar to brood sorting in ants.

Both Beckers *et al.* and Melhuish *et al.* focused on creating minimalistic robots, with

minimal sensors, minimal behaviours and minimal intelligence, still able to produce an emergent result such as clustering or sorting of an object. Both researchers also created experiments in which there would be no explicit inter-robot communication, but instead just stigmergic communication.

Stigmergy can be defined as essentially the production of a certain behaviour in agents as a consequence of the effect produced in the local environment by previous behaviour [10]. Both Beckers *et al.* and Melhuish *et al.* relied on stigmergy for their robots' respective tasks to be completed. Thus, the emergent behaviours of their agents were clustering and annular sorting.

Knowing that neither of these two specific experiments engaged in direct communication between robots but rather relied on stigmergic interaction, it is a logical step to continue this discussion with regard to the research that has been done in multi-robot communication.

3.2.2 Work Done in Multi-Robot Communications

Few papers have tackled the direct inter-robot communication issues in collective robotics systems. However, the researchers who have done some work have all come up with fairly similar conclusions.

Balch and Arkin [6] conclude that communication improves performance significantly in tasks with little environmental communication, and that complex communication strategies offer little or no benefit over low-level communication. Dahl *et al.* [15] concluded that the individual controllers using communication are significantly better than those if the controller is communication-free. Easton and Martinoli [18] also find that communication schemes heavily influence the rate of successful collaborations. Finally, Matarić [43] concluded that communication may compensate for the limitations of more direct sensory modalities, thus enabling learning in multi-robot systems. More specifically Matarić discussed the use of communication to reduce the undesirable effects of locality in fully distributed multi-agent systems with multiple agents learning in parallel while interacting with each other [43].

In Stone and Veloso's survey of multi-agent systems [5], they discuss the effects of communication in heterogeneous and homogeneous multi-agent systems. Several examples are given including [26] one in which several heterogeneous robots use their different capabilities to perform a mapping and exploration task, with communication existing between the lower level robots and the main "leader" robot and not specifically between all agents at any time. It is the team leader's task to fuse the data together and propagate it back to the "lesser" robots. Also based on Stone and Veloso's survey, it appears that not many homogeneous systems with communication have been developed with the exception of the state/goal communication systems [6] that we previously mentioned.

All of these researchers tested the effects of one or more different types of communication

on a multi-robot system. All of them generally concluded that adding communication to the system provided a better result based on their individual performance metrics. However, after completing our literature review we still found that some aspects of communication and perception could be further pursued.

3.2.3 What Has Not Been Done?

Our literature review did not find any evidence of physical robots that were able to perform the segregation sorting task. We also noticed that after reading most of the literature about multi-robot communication, there had been no attempts to establish any relationships between perception and communication. Can communicated data offset any lack of perception/sensing ability? Does enhanced perception/sensing ability offset inter-robot communications?

3.3 Why?

As explained in Sections 2.2.3 and 2.3.2 there are several reasons for developing behaviour-based robots and collective teams of robots. We have also noticed that systems capable of sensing great amounts of information are both expensive to buy and have inherent operational expenses in the hardware used to process the sensed data. We believe that having fewer sensing abilities (e.g., a smaller camera), thus a smaller realm of perception, and adding a layer of communication can be equivalent to a system with greater sensing abilities. Thus, all that is needed is a communication protocol and an understanding of what will be communicated rather than the expensive sensing system. Another example of using multiple robots and communication to more effectively accomplish team goals arises when terrain provides obstacles that can cause a robot to be blind to certain data (e.g. a cliff or hill). However, via communication, a robot can "see" behind obstacles and thus make better decisions with more information available to it.

3.4 Thesis Statement

We propose that in multi-robot systems there is a relationship between perception and communication. We propose that the addition of simple inter-robot communication to robots with lesser sensing abilities can offset these deficiencies in local sensing.

3.5 Summary

In this chapter we explained why we chose the sorting task as our problem environment. We then detailed what has been specifically done in both the collective sorting research area and the multi-robot communications research area. We concluded with why our research is

interesting and our thesis statement. The next chapter details some of the theory required to accomplish our goals.

Chapter 4

Theory

4.1 Introduction

In this chapter we present the theoretical background necessary to understand this thesis. As mentioned in Chapter 3 and further specified in Section 3.4, the goal of this thesis is to examine the correlation between perception and communication in a multi-robot environment. To accomplish this goal, we chose to implement the task of sorting since it was a task that had been experimentally proven [46, 10]. Next, a framework for the experiments must be presented and clarified. This section aims to provide theoretical background with respect to the various areas of the experiments that were performed, which in turn required us to define perception, communication, sorting, density and efficiency. The following sections define these terms in relation to the experiments.

4.2 Perception

In the robotics community the term perception can imply many different things. Webster [47] defines perception as an awareness of the environment through physical sensation. Perception, is generally acknowledged to be a representation of what is perceived. However, it is the question of what is perceived, that changes many opinions on what perception is. Perception can also be thought of as the basic component in the formation of a concept. Finally, perception can be the translation of data from a transducer into information that is useful to the operator of the information.

Of course before anything can be perceived it must be sensed and thus a brief definition of sensing is listed below. In this thesis sensing is defined as the physical stimulus that a robot receives via its sensors. Essentially, sensing means nothing to a robot; it is the interpretation of the information from the sensors that gives any sort of meaning to the term sensing. Thus, it is this interpretation of the information that yields perception.

4.2.1 Necessity/Relevance

Perception is the result of processing the sensed data into a form meaningful to the robot. This data is used to help the robot make its next decision. Throughout this document the term perception will refer to the objects that have been sensed and identified by a specific robot. I.e., the robot's perception includes all objects and other robots that it can sense (or see) and their geographic location with respect to the robot.

Similarly the term sensing will be defined as the area range that the robot is able to sense (or see). In this area, there may be objects, robots, blank space, or other world items. However, once the sensed data has been passed into the perception space, only the data that has meaning to the robot will be processed and thus interpreted.

4.2.2 Definition

The decision making process a robot undertakes is stimulated almost entirely by what the robot can perceive or "see." Before any action or behaviour can be performed the robot must be able to convert the sensory data into useful information. As demonstrated in Figure 2.2 actuation only occurs after the sensory data enters the system and triggers specific behaviours.

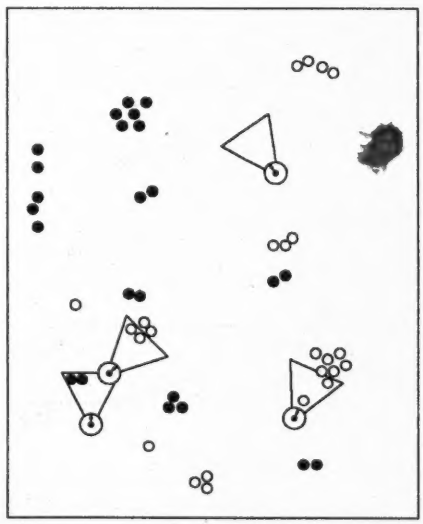

Figure 4.1: A sample world, W. The white and black circles are objects of type 1 and 2. The larger white circles with black circles inside are robots and the arrows in the circle signify their orientation. The triangular areas coming from a robot is its area that is sensed s_r.

The following paragraphs define the world, the sensing area of a robot, the objects a robot perceives and the total perceived space of a robot. Please refer to Figure 4.1 to get a better picture of what the robot's world may be like.

Before defining perception, several symbols need to be explained. In Equation 4.2 x_r and y_r are the coordinates of robot r in W, θ_r is the orientation of r in W, δ_r is r's sensing distance and ω_r is r's sensing angle.

In Equation 4.1 v_k is defined as a vector consisting of 6 members. These members include an x_k and y_k coordinate signifying the center of gravity of the object identified, three booleans denoting if the object is of type 1, type 2 or a robot (p_k^1, p_k^2, and r_k) and θ_k, which represents the orientation of the object identified. However, θ_k is only applicable when the object identified is a robot. It is also necessary to note that p_k^1, p_k^2 and r_k are of type $\mathbf{B} = \{0..1\}$ and that only one of p_k^1, p_k^2 or r_k can be equal to 1 for any given object.

$$^W v_k = [x_k, y_k, p_k^1, p_k^2, r_k, \theta_k] \tag{4.1}$$

Finally Equation 4.2 defines the perceived information known to robot r. In Equation 4.2 $\mathbf{N}_t$ refers to the number of visible objects of type t and is defined as $\mathbf{N}_t = \{0..n_t\}$ where n_t $= n_1$ or n_2 is the total number of type-t objects in the experiment ($t = 1$ for type 1 objects, and 2 for type 2 objects). Similarly, $\mathbf{N}_r$ refers to the number of visible robots and is defined as $\mathbf{N}_r = \{0..n_r\}$ where n_r is defined as the total number of robots in the experiment (1, 2 or 4).

$$\mathcal{P}_k = f(x_r, y_r, \theta_r, \delta_r, \omega_r) = [v_1, v_2, ..., v_\kappa] \quad , \quad \kappa = \mathbf{N}_1 + \mathbf{N}_2 + \mathbf{N}_r \tag{4.2}$$

It can be concluded that perception $\mathcal{P}_k$ is the mapping of sensory world data from the robot's sensors into a perception space that can be interpreted by the robot.

4.3 Communication

Communication has many different meanings in the robotics community. Communication can be explicit or implicit. I.e., communication can be thought of as direct either robot to robot communication (explicit) or as observed robot interactions on the world (implicit or stigmergic). Explicit communication can be as simple as observed binary data regarding the state a robot is in or as complex as exact GPS coordinates and visual data seen by a robot. Implicit communication is more difficult to discretise since all implicit communications depend on how well a robot can perceive the world and how the robot's interactions have affected the world.

4.3.1 Necessity/Relevance

Communication is used both implicitly and explicitly in the experiments outlined in Chapter 5. An entire set of experiments without explicit communications were run as well as an entire set of experiments with explicit communications.

The experiments that were not run with explicit communication still involved implicit or stigmergic communication. Robots were able to perceive different classes of objects and robots. If an object was too close to a robot (the center of the object was within 6 cm of the edge of a robot), then the robot implicitly understood that that particular object was occupied. However, even though a robot could identify when another object was a robot, it could only "see" that robot strictly as an obstacle in the stigmergic situation.

The experiments that were run with explicit communication had the same implicit communication described above. However, when a robot saw another robot, it also got explicit information back from that robot. This explicit information conveyed was precisely what the second robot perceived at that instant in time. Figure 4.2 shows in more detail what explicit information was conveyed.

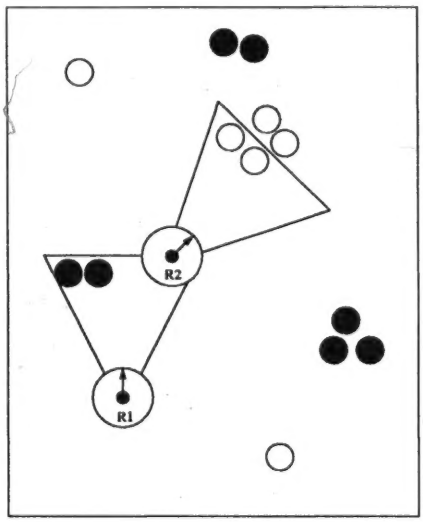

Figure 4.2: Communicating robots (R1 can perceive what R2 can perceive).

4.3.2 Definition

As explained in the previous section explicit information was communicated from one robot to another if it was "seen" by another robot (Figure 4.2). Equation 4.4, shows that the communicated information is simply a perceived geographical area that the robot can use to make its decision. Also, shown in Equation 4.3, you can see that the overall perceived data understood by a robot that has received communicated information is simply the union between the perception of robot i and other robots j that robot i can "see."

$$\mathcal{P}'_i = \mathcal{P}_i \bigcup_{j=1}^{N_r} \mathcal{P}_j, \quad \forall \quad {}^{W}v_i \quad \exists \quad r_i \neq 0 \tag{4.3}$$

So just exactly what information is being communicated? As can be seen in Figure 4.2 the communicated information is simply all of the visible robots' perceived information. This is further demonstrated in Equation 4.4 where the total communicated information received by robot i is the perceived information $\mathcal{P}_i$ subtracted from the new perceived information $\mathcal{P}'_i$.

$$\mathcal{C}_i = \mathcal{P}'_i - \mathcal{P}_i \tag{4.4}$$

Equation 4.4 can perhaps be better understood by looking at Equation 4.5 which elaborates on the example shown in Figure 4.2. In Equations 4.5 and 4.6 the communicated information from robot 2 to robot 1 is labeled $C_{2,1}$. In general $C_{2,1} \neq C_{1,2}$.

$$\mathcal{C}_{2,1} = \mathcal{P}'_1 - \mathcal{P}_1 = \mathcal{P}_2 - \mathcal{P}_1 \cap \mathcal{P}_2 \tag{4.5}$$

and combining the equations from Equation 4.5 we get

$$\mathcal{P}'_1 = \mathcal{P}_1 \cup \mathcal{P}_2 = \mathcal{P}_1 \cup \mathcal{C}_{2,1} \tag{4.6}$$

It can be concluded from this section that total communication $\mathcal{C}_i$ to robot r_i is simply the perceived information of robot r_j minus any previously known information of r_i via communication or perception.

4.4 Sorting Method

In Section 2.3.3 several examples of sorting and foraging techniques are given. In general the foraging methods that have been used in the past have had robots operate using two basic rules: pick-up and put down. Sometimes the robots brought objects back to a central location and other times clusters just formed in miscellaneous places until eventually a larger cluster emerged.

As explained in the the Figures 2.3, 2.4, 2.5 and 2.6, there are 4 different types of sorting. The clustering figure (Figure 2.3) explains the clustering/foraging version of sorting while the other three figures explain types of sorting multiple objects from each other. There are very basic pick up and drop rules used in all of these different methods, similar to those used by the clustering method.

4.4.1 Necessity/Relevance

All of the experiments in this thesis implement the segregation sorting method. The robot collective is homogeneous and each robot employs an identical sorting algorithm. A robot can be in one of three states based on its position in the world and where its target objects are. These states are: "wander," "pick up object," and "drop object." With no objects

visible within its perceptual space, a robot is in the wander state and will execute a random walk. Once one or more objects appear in its perceptual space, if the robot does not hold an object already and the pick up rule has been evaluated as true, the robot moves to the pick-up state where it will select an object to pick up. When the robot acquires an object, it moves to the drop state to decide where and when to put down the object.

4.4.2 Definition

The sorting algorithm used in this research was developed using a very simple state machine which involved a few rules. The state machine is shown in Figure 4.3 and the rules are explained in the following sections.

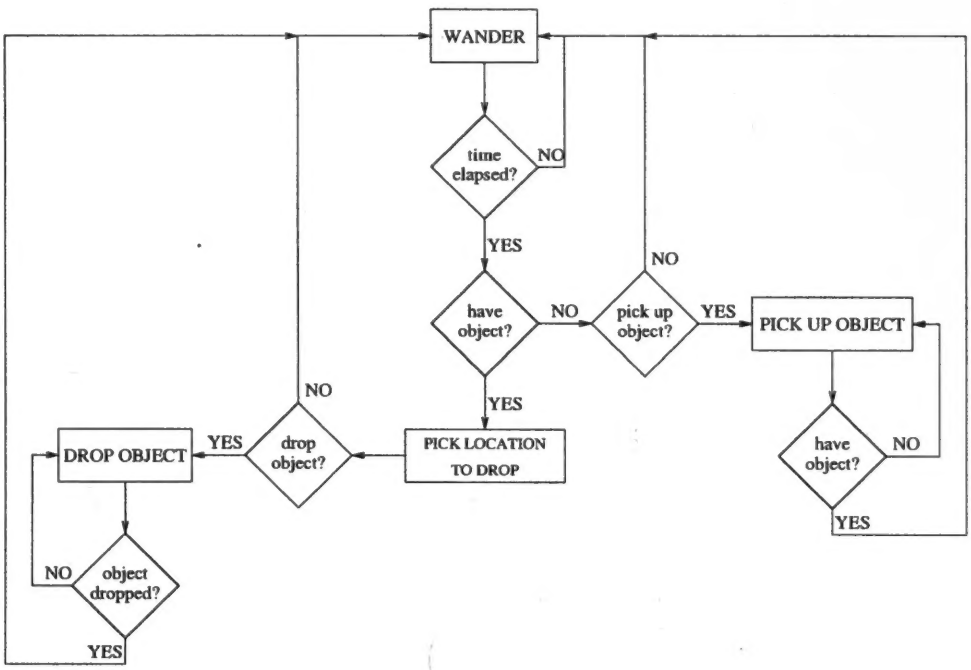

Figure 4.3: Sorting algorithm state machine

Pick Up Rule The pick-up rule allows a robot to decide if it should approach an object or which object to pick up within its perceptual space P_i'. It is a function of both the distances to the closest objects of the two classes, the densities of the objects of the two classes, and the numbers of the objects within P_i'. $Dist_1$ and $Dist_2$ is defined to be the distances from the robot to the closest type 1 and type 2 objects, respectively. The robot will pick up the object with the smaller ratio determined by Equation 4.7.

$$R_i = \frac{Dist_i}{\mathcal{U}_i},\ i = 1, 2 \tag{4.7}$$

$\mathcal{U}_i = f(\rho_i,\ m_1,\ m_2)$ is defined in Equations 4.8 and 4.9 where ρ_i and m_i are the density and the absolute number of the objects of type i in the perceptual space for the given robot. The function $f()$ represents a composite heuristic measurement of the sparsity of objects of a particular type, and its details can be found in Equations 4.8 and 4.9.

$$\mathcal{U}_1() = \begin{cases} \frac{1}{4}(1 - \frac{\rho_1}{\rho_1^*}) + \frac{3}{4}(1 - \frac{m_1 - m_2}{\alpha}) & : \quad m_1 > m_2 \\ \frac{1}{4}(1 - \frac{\rho_1}{\rho_1^*}) + \frac{3}{4} & : \quad m_1 <= m_2 \end{cases} \tag{4.8}$$

$$\mathcal{U}_2() = \begin{cases} \frac{1}{4}(1 - \frac{\rho_2}{\rho_2^*}) + \frac{3}{4}(1 - \frac{m_2 - m_1}{\alpha}) & : \quad m_2 > m_1 \\ \frac{1}{4}(1 - \frac{\rho_2}{\rho_2^*}) + \frac{3}{4} & : \quad m_2 <= m_1 \end{cases} \tag{4.9}$$

Once the decision to pick up an object is made, the robot approaches the object while continuously evaluating $\mathcal{U}_i$ based on the robot's surroundings to account for the perceived changes in the robot's P_i'. Through experimentation with the pick-up equations (Equations 4.8 and 4.9) it was decided that the equations would be 75% based on the number of objects in an area directly in front of the robot and 25% based on the density of the objects in that area. The definition of "density" is given in Section 4.4.3. This was done for one major reason. During experimentation, if equal weights were placed on the two variables (density and number of objects), we observed that the density was so low in the beginning stages that objects would always be picked up. This was good except small piles rarely became large piles because the objects were always being picked up and thus the segregation sorting took much longer than desired.

Drop Rule The drop rule allows a robot to decide where and when to put down an object that it is holding. The decision on whether an object should be put down is determined probabilistically, similarly to [16]. Specifically, we define a drop function $\mathcal{D}_i = g(\rho_i,\ m_1,\ m_2)$ whose value lies between 0 and 1, indicating the likelihood of dropping the object being held, i.e., the closer the value to 1, the more likely the object will be dropped. $\mathcal{D}_i$ is a heuristic function we have defined and its details can be found in Equations 4.10 and 4.11. Once the decision to drop the object is made, the question of where to put it down, assuming a type-i object, is resolved by first calculating the center of gravity of the objects of type i within the perceptual space of the robot, and then locating the object closest to the center of gravity. This closest object is used as the reference point for the drop location. Note that we continuously evaluate this location as the robot moves to reflect changes in the robot's perceptual space.

$$\mathcal{D}_1() = \begin{cases} \frac{1}{2}(\frac{\rho_1}{\rho_1^*}) + \frac{1}{2}(\frac{m_1 - m_2}{m_1}) & : \quad m_1 > m_2 \\ \frac{1}{2}(\frac{\rho_1}{\rho_1^*}) & : \quad m_1 <= m_2 \end{cases} \tag{4.10}$$

$$\mathcal{D}_2() = \begin{cases} \frac{1}{2}(\frac{\rho_2}{\rho_2^*}) + \frac{1}{2}(\frac{m_2 - m_1}{m_2}) & : \quad m_2 > m_1 \\ \frac{1}{2}(\frac{\rho_2}{\rho_2^*}) & : \quad m_2 <= m_1 \end{cases} \tag{4.11}$$

Through experimentation with the drop equations (Equations 4.10 and 4.11) it was decided that the rule would be 50% based on the number of objects in an area directly in front of the robot and 50% based on the density of the objects in that area. We initially started with the 50/50 ratio and after experimentation we observed that objects were being dropped in appropriate situations and we left the ratio intact.

4.4.3 Maximum Density

Hales' honeycomb conjecture [29] offers proof that any partition of a plane into regions of equal area has a perimeter at least that of the regular hexagonal honeycomb tiling (Figure 4.4).

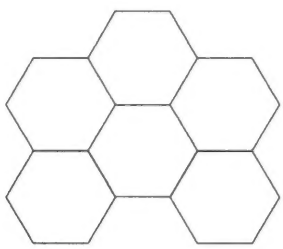

Figure 4.4: Honeycomb pattern

We have expanded upon Hales's proof and have used ρ_t^* to define the maximum density that a group of circular objects can have where t is simply the type of object for which we are finding the maximum density. In Equation 4.12 d_T is the summed distance of m objects to c_g, the center of gravity of all the objects.

$$\rho_t^* = \frac{m_t}{d_T} \tag{4.12}$$

$$d_T = \sum_{i=1}^{m} d = \sum_{i=1}^{m} \sqrt{(x_i - x_{c_g})^2 + (y_i - y_{c_g})^2} \tag{4.13}$$

Density and maximum density are calculated using the same formulas. The only difference is that d_T in the optimum case is much smaller than in the regular case and when ratios of densities are taken for the pick-up and drop rules (Equations 4.8, 4.9, 4.10 and 4.11) the ratio is closer to 1 when the objects are more compact.

4.5 Efficiency Calculation

The efficiency of our system is based on the percentage complete versus the time to complete the desired task. In order to define efficiency we first need to define how we will calculate percentage complete.

4.5.1 Definition

Based on Equation 4.14 we define efficiency as the completion rate over time, i.e. the faster a task is completed with respect to time the more efficient the system.

Equation 4.14 defines the "%complete" where n is the number of piles or clusters of Type-A objects, and each cluster contains A_i objects.

$$\%complete_A = 100 \times \sum_{i=1}^{n} \left(\frac{A_i}{A_{total}} \right)^2 \tag{4.14}$$

This percent completion will be calculated for all types of objects and the percent completion values will be averaged to arrive at the overall system percent completion measure. Figure 4.5 details just exactly how the percent complete equation is performed.

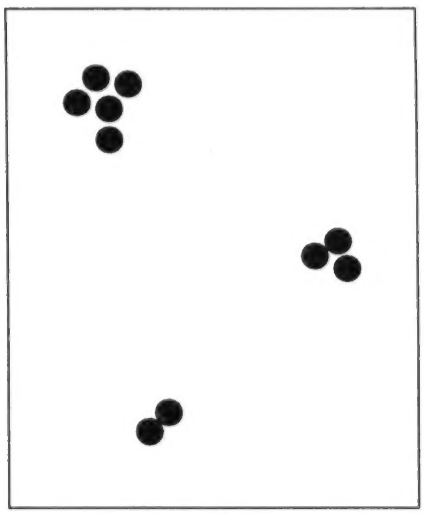

Figure 4.5: This figure illustrates Equation 4.14. Given one pile with three objects, one pile with five objects and the last pile with two objects, the percent complete is $(\frac{3}{10})^2 + (\frac{5}{10})^2 + (\frac{2}{10})^2 = 38\%$ complete.

4.6 Summary

In this chapter a framework for the experiments was presented, along with the theoretical background with respect to the various aspects of our experiments. This included definitions of perception, communication, sorting, density and efficiency. The following chapters use these definitions to describe the set-up and execution of the experiments.

Chapter 5

Experimental Setup

5.1 Introduction

To experimentally study the collective sorting problem and evaluate our segregation sorting algorithm under varying parameters we designed a set of experiments using RoboCup Small-Size League robots [58]. The experimental environment used was also very similar to the environment used by RoboCup Small-Size League robots.

The entire experimental system was created for two purposes. The first was to enhance multi-robotics research at the University of Alberta. The second was to create a platform and system in which a Small-Size League RoboCup soccer team could be developed. The complete system was comprised of the robots, a vision system, a wireless communications link, robot monitoring tools and overall system integration.

Once the system was developed for RoboCup it was slightly modified for multi-robot research purposes. In the following sections, we describe some of the assumptions we made in the experimental process, we detail the robots hardware/firmware, we illustrate the experimental environment, we describe the vision system, we explore the robot monitoring system, and finally, we recite the methods used to perform our experiments.

5.2 Assumptions

The assumptions we made while conducting our experiments are listed below:

- It was assumed that an object within 2 cm of the robot's front is being controlled by that robot. This assumption conveyed the existence of a form of stigmergic communication.
- We assumed that the robot knew how to stay away from the 5 cm walls that bound the environment. The software agents controlling the robots ensured that robots stayed 10 cm away from the walls so that the grippers on the robots would not get caught on the walls.

5.3 Robots

The robots were built to the standard RoboCup Small-Size League rules [58] and modified slightly for the collective sorting problem. The robots were designed and built from scratch at the University of Alberta. In the following subsections we will discuss the mechanical, electrical, and firmware specifications.

5.3.1 Mechanical

Each robot consisted of two horizontal metal plates one on the top and one on the bottom. The two metal plates were displaced by 12 cm and were joined to each other with four standoffs placed symmetrically around the plates. The plates were 18 cm in diameter with a 10 cm chord cut out off front. Everything on the robot was attached to one of these two metal plates.

Most of the electronics for the robot as well as the robot identification top were attached to the top plate. The identification top consisted of a blue dot and three pink or green dots on the front edge. The blue dot was used to identify the center of the robot and the pink and green dots were used as identification markers. There were three printed circuit boards attached to the top plate. All the electronics were connected to the top plate via metal and plastic standoffs.

The robot's drive train, situated on the bottom plate, consisted of two motors each placed strategically in the exact center of the robot. Essentially, the robot could turn in place and it's center blue dot would not translate any distance. A small section of electronics was also attached to the bottom plate. This circuit board was used to control IR (InfraRed) LEDs (Light Emitting Diodes) on the front of the robot, which were used to signal if a robot had an object or not. The grabber was also connected to the bottom plate and it consisted simply of two non-moving curved arms that were able to corral an object and not lose it unless the robot translated directly backwards. The grabber was static and didn't move. An image of the grabber attached to the front of the robot is shown in Figure 5.8.

5.3.2 Electrical

Four circuit boards were used in the robot. Two were off-the-shelf products, one was a printed circuit board made specifically for motor control for the robot, and one was a small circuit board created for the IR beam breaking circuit. Along with the circuit boards, other electrical components were used including motors and standard IR LEDs.

The main controller board was designed and manufactured by RoboMinds. A picture of the board is shown in Figure 5.1. It is identified as a Mini RoboMind (MRM). The MRM contains a MC68332 microprocessor, 512K of Flash, 512K of RAM and measures 2.95 inches square [59]. The MRM is easily programmed in either assembler or C programming

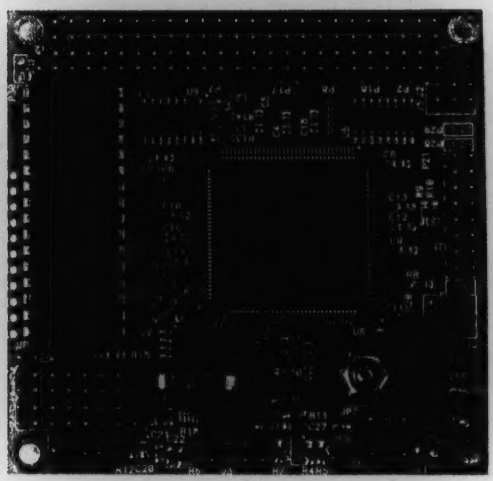

Figure 5.1: The main controller board used in our robots. The controller is based around the MC68332. It contains 512K of Flash, 512K of RAM and measures 2.95 inches square.

languages. Even though the MC68332 offers several features only a few were utilized for this robot. The Multichannel PWM (Pulse-Width Modulation) TPU Function (MCPWM) routines were run on TPU (Time Processing Unit) channels 0-2. These channels were used to operate the two driving motors. The MCPWM function uses two externally gated TPU channels to produce sophisticated PWM signals. These PWM signals are used with the H-Bridges to control the speed of the motors. The two TPU channels that are used are externally gated using an XOR gate. Each channel generates a 50% duty cycle PWM signal. The resulting PWM signal is modified by changing the phase difference between the two channels. Figure 5.2 show examples of how the two TPU channels are used to produce the desired PWM signal.

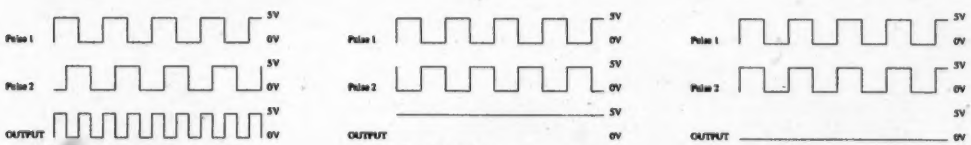

(a) 50% duty cycle (b) 100% duty cycle (c) 0% duty cycle

Figure 5.2: An example of how various PWM duty cycles are produced by XORing two phase-shifted pulse trains using the MC68332 MCPWM routines.

Complex motion control is achieved by adding feedback from the motors into a control loop to control the motor's speed. The first thing that is necessary is to obtain the feedback information from the motors. The two motors have two 16-bit feedback lines that can be used to determine the distance a motor has turned. The two signals can also be used together to determine the direction that the motor is turning. The feedback information

from the motors is interpreted using the Fast Quadrature Decode (FQD).

The FQD TPU input function uses two channels to decode a pair of out-of-phase signals in order to increment or decrement a counter. It is particularly useful for decoding position and direction information from a slotted encoder in motion control systems. This function uses a pair of adjacent TPU channels to decode quadrature signals into a 16-bit counter located in PRAM. The counter is updated when a valid transition is detected on either one of the two inputs. The user can read or write to the counter at any time. The counter is free running overflowing to 0x0000 or underflowing to 0xFFFF depending on the direction. The FQD routines were run on TPU channels 12-15 and were used along with the MCPWM routines to gather velocity information returning from the drive motor encoders.

The Queued Output Match TPU Function (QOM) routines were run on TPU channels 4-7. These channels were used to set the direction of each of the motors to forward or reverse. Finally, PORT E was set up to handle all of the wireless radio communications.

The wireless communications board was designed and manufactured by Radiometrix. A picture of the board is shown in Figure 5.3. It is identified as an RPC-418-40 (or RPC-433-40).

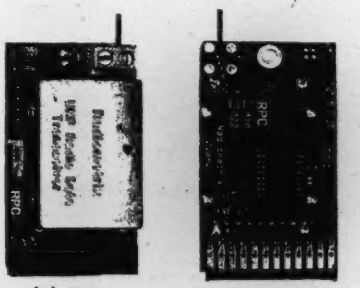

(a) front view (b) back view

Figure 5.3: The wireless communications board we used in our robots. The RPC-418-40 is designed and manufactured by Radiometrix.

The RPC-418-40 is an intelligent transceiver module which enables a radio network/link to be easily implemented between a number of digital devices. The module combines a UHF radio transceiver and a 40kbit/s packet controller [56]. The RPC-418-40 is a self-contained plug-on radio port which requires only a simple antenna, 5V power supply and a byte-wide I/O port on a host micro-controller. The RPC-40-18 provides all the RF circuits and processor intensive low level packet formatting and packet recovery functions required to interconnect any number of micro-controllers in a radio network [56].

The motor controller board was designed specifically for these robots. An example is shown in Figure 5.4. This board was designed using the L298N as the high voltage high current dual full-bridge driver. Three of these chips were used on this board, although only

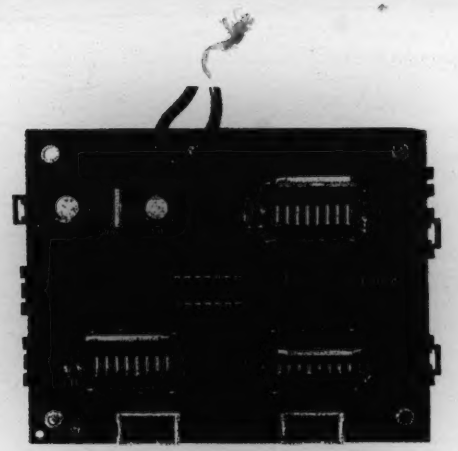

Figure 5.4: The motor board used in our robots. This board contains all the high current components and controller components necessary to drive a robot's motors.

two were used for the robot's motor drivers in these experiments. In addition to the L298N chip, a simple XOR chip was used on the board to create and relay the proper signals for the motors from the MCPWM routines discussed above and illustrated in Figure 5.2.

The IR beam breaking circuit board was designed simply to detect when an IR beam had been broken. A picture of the board is shown in Figure 5.5.

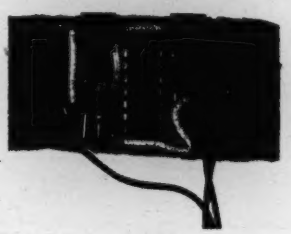

Figure 5.5: The IR break beam board used in our robots. This board simply two voltages and returns a 1 if the compared voltage is higher than the reference voltage or a 0 if the compared voltage is lower than the reference voltage.

The IR board compares two voltages using the LM339 quad comparator (although only one of the comparators on the chip was used). If the compared voltage is higher than the reference voltage a 1 is returned, otherwise a 0 is returned. The 1 signifies that the beam is broken and the 0 signifies that the beam is not broken.

The final major electrical components are the motors we used on our robots. A picture of the motor is shown in Figure 5.6. It is identified as the MicroMo 2224U006SRIE2-16+20/1 23:1+X0758C. The motors for the drive train system are capable of a torque up to 0.71 oz-in; with the 23:1 gear ratio we used, this torque becomes 16.3 oz-in.

Figure 5.6: A robot motor. This motor is able to deliver full quadrature decoding signals.

5.3.3 Firmware

The firmware on the MRM had two simple roles. The firmware's first task was to use the MCPWM and FQD functions of the MC68332 along with a Proportional Integral Differential (PID) Controller to control the velocities of the output motors. Its second function was to manage communications with the base station.

We were able to develop a reasonable model of the PID controller by varying each of the K_p, K_i, and K_d constants and monitoring the output stability. K_p, K_i, and K_d, signifies the PID controllers' proportionals, integral and differential gain respectively. Initially the K_i and K_d constants were set to 0 and the system was tuned to a relatively stable state by varying the K_p constant. Using only K_p to tune the model we still experienced some oscillations around our specified velocity. To tune this with the K_i constant still zero, we varied the K_d constant to dampen the response and further improve on the controller's stability. After achieving a stable system, we attempted to add a K_i constant into our controller to decrease the response time even further. However, after experimenting with the robot and the constants we found that adjusting the K_i constant provided negligible benefit to the system and accordingly we decided to use a PD controller to control the motor velocity. This controller is shown in Figure 5.7.

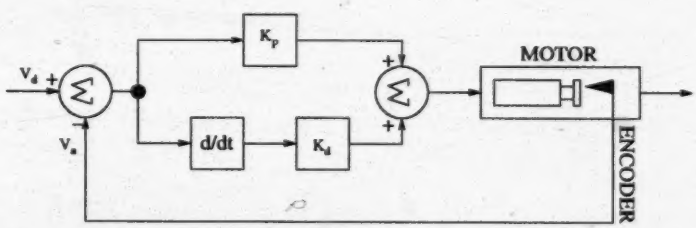

Figure 5.7: The PD controller used to control the motor velocities. The K_p and K_d constants were tuned experimentally. V_a is the actual velocity as sensed by the motor encoder and V_d is the desired velocity specified to the controller.

In order to manage communications with the base station, a very basic controller was designed for the robot based on the specifications from Radiometrix. All details about the packet format can be found on the Radiometrix website [56].

5.3.4 Complete Robot

The complete robot consists of the metal plates, the electronics, the firmware, the motors, the wireless radio and the grabber. A picture of one of the robots is shown in Figure 5.8.

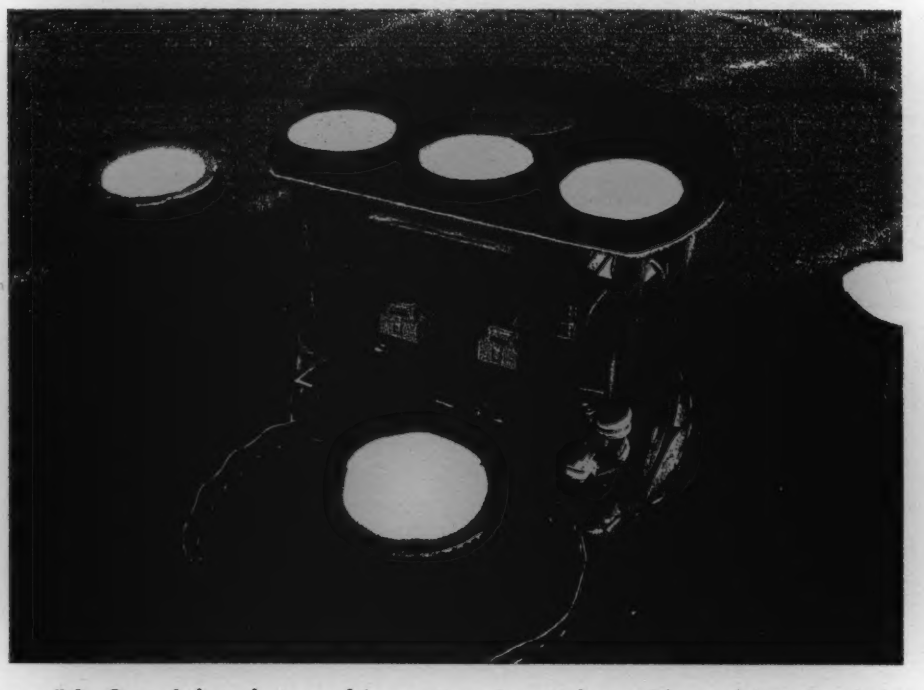

Figure 5.8: One of the robots used in our experiment along with an object that is being controlled by the robot.

5.4 Environment

The experimental environment we used as the robots' "world" was rectangular in shape, had a carpeted surface, measured $2m \times 3m$ and was surrounded by a 5 cm high wall. An overhead camera identified the robots and collected visual data of the entire surface. This data was processed and then used by software agents that ran on isolated computers and communicated with the robots via a wireless communication link.

5.4.1 Vision System

Our vision system used the IEEE DragonFly camera from Point Grey Research Inc. The Dragonfly is a compact, fully digital IEEE-1394 board level camera. The Dragonfly uses a 1/3" progressive scan CCD in order to stream VGA quality color images at 30 FPS without compression [28]. We used it to sample 640x480 24 bit RGB color images at 30 fps into a

normal PC with an off-the-shelf Firewire capture card, an AMD Athlon XP 1700+ processor and 512 MB of RAM [45].

This camera identified the robots and collected visual data of the entire surface A complete explanation of how the vision system worked can be found in [45]. This data was then parsed to provide the robot agents with only information about what they could actually "see" based on their heading and desired sensing and communication parameters.

5.4.2 Wireless Communications

As described in Section 5.3.2 we used the Radiometrix RPC-418-40 for our wireless communications. The Eval-RPC (PIC16F84A-20I/P micro-controller) was used as a base station that passed data to a PC, which in turn conveyed messages to all of the robots. The Eval-RPC is based upon a modified RPC-000-DIL design and provides management of all the necessary data and control lines. Figure 5.9 is a photo of the base station.

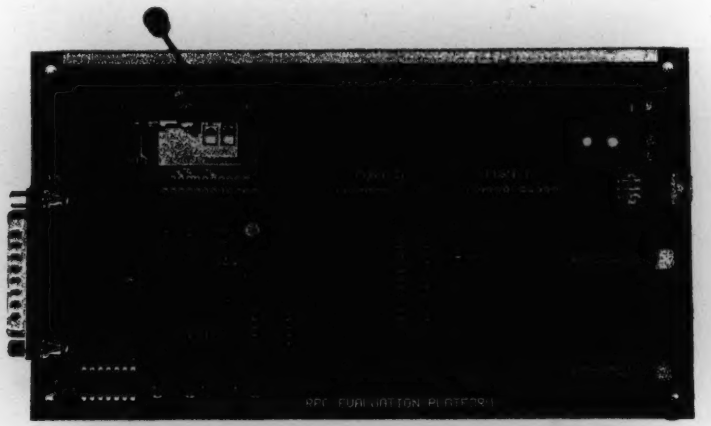

Figure 5.9: This is the base station used to transmit signals to and receive signals from the robots.

5.4.3 Robot Monitoring

A robot monitoring system was created to monitor all of the robots and objects in the environment. This system kept track of robot movements and enabled a single click of a mouse button to start and stop the experiments according to the completion criteria outlined in Section 4.5. A screen shot of the robot monitoring system is shown in Figure 5.10.

5.5 Methods

At the beginning of each of the experiments, n robots were placed in the center of the area along with 24 Type-1 objects (orange) and 24 Type-2 objects (yellow) in a perfectly

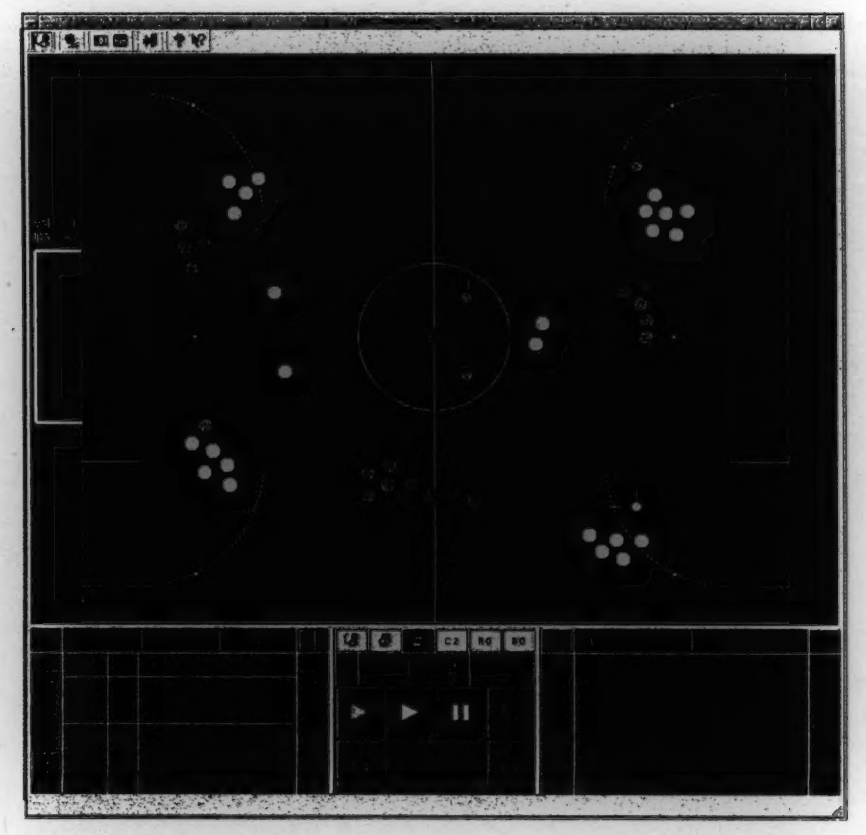

Figure 5.10: A screen shot of BotCommander. BotCommander (the robot monitor) was used to monitor the robots and objects in the experiments. The experiments could be started or stopped with a click of the mouse.

mixed symmetric pattern within the sorting area (Figure 5.11). This assured the same initial condition for all the experiments. The robots were inspected and the batteries were recharged before each set of experiments. The agents were run and data was collected at 30 second intervals throughout the experiment.

Several sets of experiments were run varying different sets of parameters. The number of robots per experiment varied from one, two or four robots. The robot's sensing ranges and communication ranges were set at 50 cm, 100 cm or 180 cm respectively. Finally, communication was either available or not available during the experiments. This allowed $3 \times 3 \times 2$ experiments to be run, with the exception of the 1 robot case, where communication was irrelevant because there were no other robots to communicate with. Therefore a total of fifteen different sets of experiments were run. All of the experiments ran until the objects were $\geq 90\%$ sorted (as defined in Equation 4.14) and each experiment was run twice.

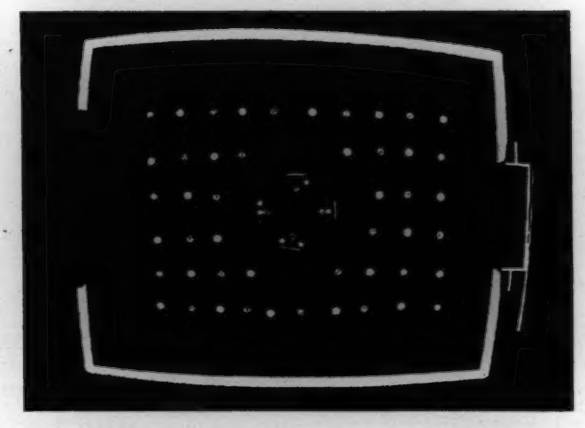

Figure 5.11: Initial arrangement of the objects (24 objects for each class), which are uniformly distributed and perfectly mixed. Note: There is considerable radial distortion in the camera causing the rectangular field to appear to be barrel-shaped. Distortion correction was performed before the coordinates of the objects were calculated.

5.6 Summary

In this chapter we discussed the experimental environment for our experiments. This included a description of the assumptions we made in the experimental process, a detailing of the robots hardware/firmware, an illustration of the experimental environment, a description of the vision system, an exploration of the robot monitoring system, and finally, a recital of the methods used to perform our experiments. The results of the experiments performed in this environment are give in the next chapter.

Chapter 6

Experimental Results

6.1 Introduction

In this chapter we present the experimental results for both the single-robot experiments and the multi-robot experiments. The single-robot experiments will be presented first to prove the viability of the algorithm shown back in Chapter 4. The multi-robot experiments are then presented demonstrating the improvement on overall system performance when more than one robot is present. All discussion about the results may be found in Chapter 7.

6.2 Single-Robot Experiments

We performed the single-robot sorting experiments to prove that our algorithm could complete segregation sorting with a single robot. More importantly, the results from the single-robot experiments served as a baseline for multi-robot experiments. Three different single-robot experiments were completed, with a sensing range of 50 cm, 100 cm and 180 cm, respectively. Each experiment was allowed to run until completion (i.e. at least 90% of the objects were sorted). Note that because there was only one robot, the issue of communication was irrelevant.

For each sensing range, there were two runs of the experiments, with results summarized in the first two rows of Table 6.1. A single robot, given a 50 cm, 100 cm, and 180 cm sensing range in successive experiments, completed its task in approximately 160, 80, and 60 minutes, respectively. The experiments have shown that a larger sensing range can speed up task completion time. However, while there was a significant improvement when the sensing range increased from 50cm to 100cm, the improvement became much less pronounced when the sensing range was changed from 100cm to 180cm. Complete data for all of the single-robot experiments can be found in Section 6.3.

		without communication			with communication		
Robots	Trial	50 cm	100 cm	180 cm	50 cm	100 cm	180 cm
1	1	145	90	58	na	na	na
	2	162	77	62	na	na	na
2	1	90	55	35	73	42	40
	2	84	61	44	63	39	33
4	1	65	47	39	37	45	34
	2	55	46	46	35	37	27

Table 6.1: Time (in minutes) taken by 1, 2 or 4 robots, with and without communication present, to sort objects of two classes into two piles with sensing ranges of 50 cm, 100 cm and 180 cm.

6.3 Single-Robot Experiment Data

For each set of experiments the percentage complete was logged every thirty seconds. Equation 4.14 was evaluated at each thirty second interval and the result was logged in a data file. This data was then compiled and graphed versus time. All of the experimental data that was collected for the single-robot experiments has been graphed and is shown in Figures 6.1 6.2 and 6.3.

Figures 6.1 6.2 and 6.3 show graphs of the data collected for all of the one-robot experiments. The graphs show a plot of percentage complete over time for the various sensing ranges used. It can be seen that some of the experiments go to 100% completion while others do not. The system tried to achieve 100% completion, however, if after the completion criterion had been met the system did not go to 100% completion the experiments were stopped considering the experiment a success.

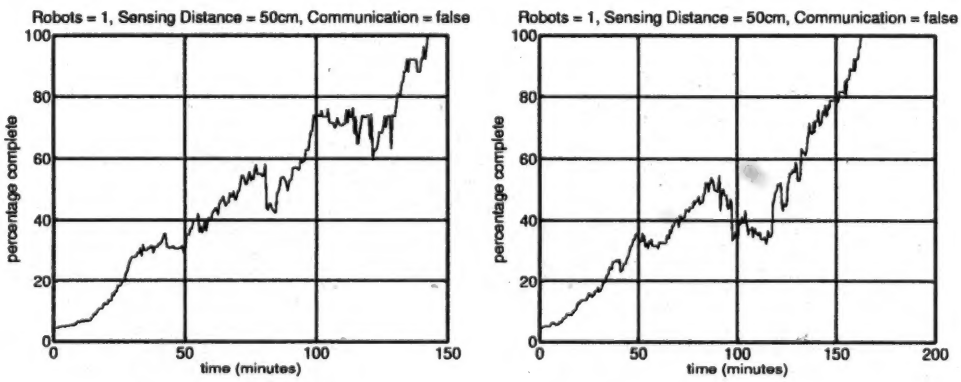
Robots = 1, Sensing Distance = 50cm, Communication = false
percentage complete
time (minutes)
Robots = 1, Sensing Distance = 50cm, Communication = false
percentage complete
time (minutes)

(a) experiment one (b) experiment two

Figure 6.1: These graphs represent the change in percent complete versus time for a single-robot experiment in which the robots has a sensing range of 50 cm.

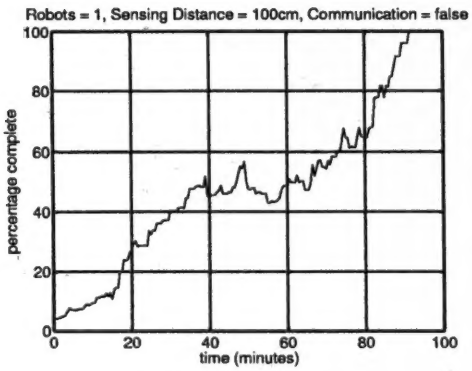

(a) experiment one

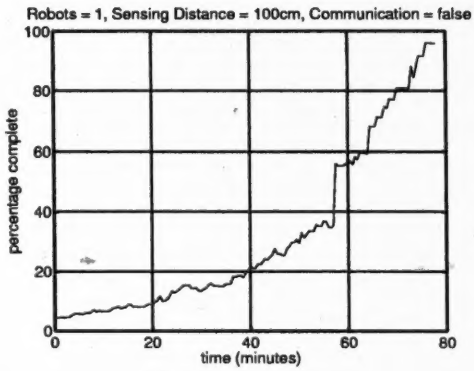

(b) experiment two

Figure 6.2: These graphs represent the change in percent complete versus time for a single-robot experiment in which the robot has a sensing range of 100 cm.

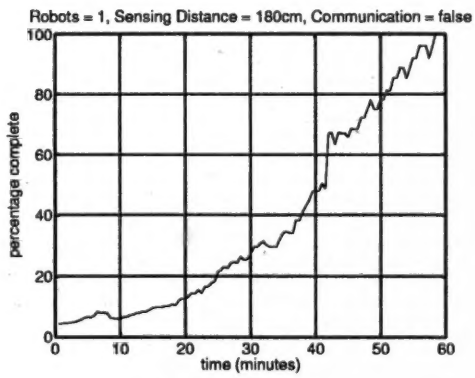

(a) experiment one

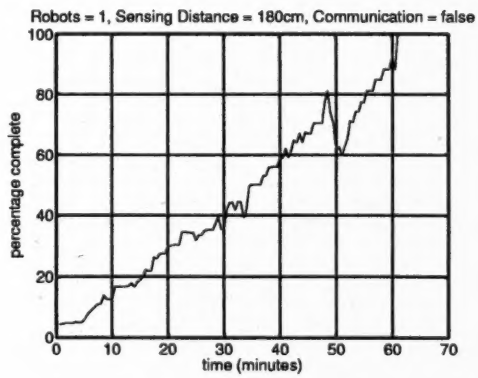

(b) experiment two

Figure 6.3: These graphs represent the change in percent complete versus time for a single-robot experiment in which the robot has a sensing range of 180 cm.

6.4 Multi-Robot Experiments

Multi-robot experiments were performed with and without communication. Three sets of system parameters were examined for both multi-robot systems. Similar to the single-robot experiments, the sensing range was set at 50 cm, 100 cm, and 180 cm for successive experiments. Each experiment was run until completion; i.e., with a percent completion value greater than 90%. For each system configuration (sensing and communication range, and the number of robots), the experiment was run twice. Figure 6.4 shows six snapshots of one experiment with four robots, at different stages of the sorting task. The results from all the experiments are summarized in the last four rows of Table 6.1. We should note that for each experiment the sensing range and communication range were equal, i.e., if a robot

could sense a robot then it could communicate with that same robot. Complete data for all of the multi-robot experiments can be found in Section 6.5.

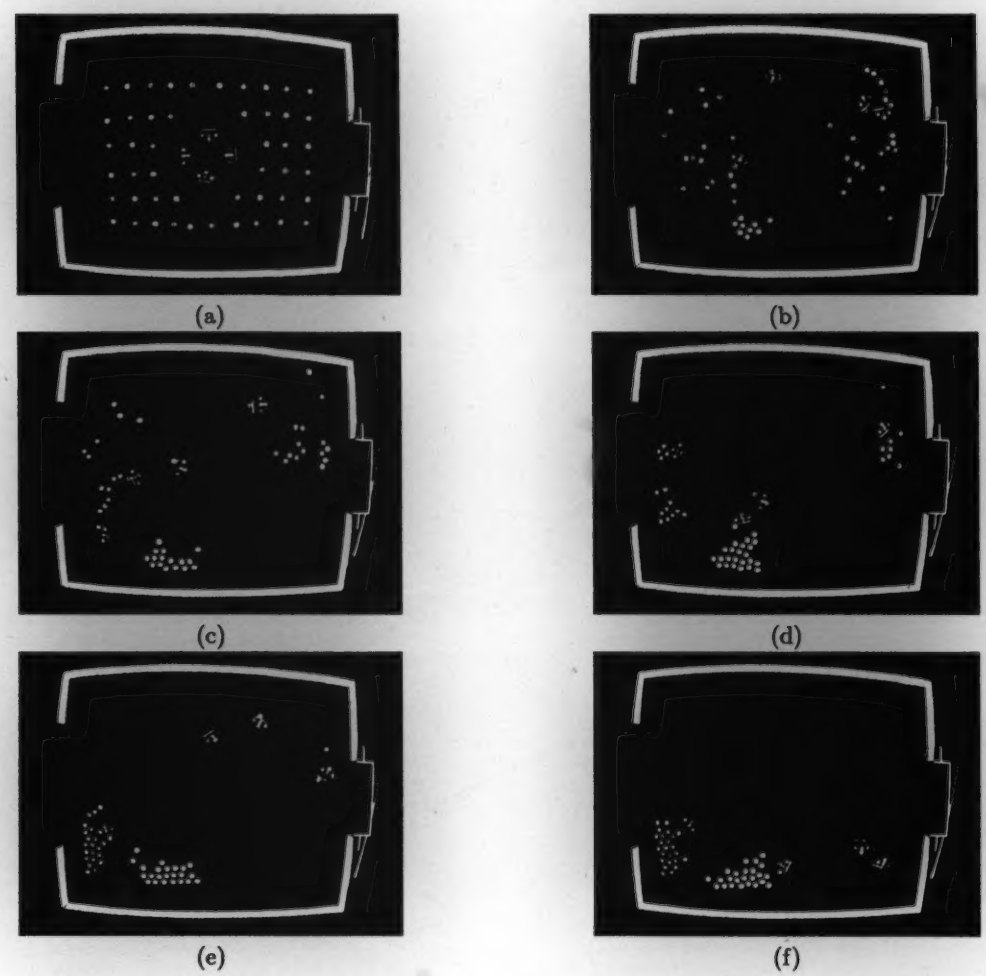

Figure 6.4: Snapshots of the experiment involving four robots, each with a 50 cm sensing and communication range. Note: There are two types of objects (yellow and orange) as in Figure 5.11, except that yellow objects have been painted white with an image editor to improve the readability of the images. (a) shows the initial conditions (b) shows the experiment status after 8 minutes, (c) shows the experiment status after 15 minutes (d) shows the experiment status after 23 minutes (e) shows the experiment status after 30 minutes (f) shows the final results of the experiment.

6.5 Multi-Robot Experimental Data

For each set of experiments, Equation 4.14 was evaluated at each thirty second interval and the result was logged in a data file. This data was then compiled and graphed versus time. In the following pages, all of the experimental data that was collected for the multi-robot

experiments was graphed and is shown in Figures 6.5 through 6.16. Table 6.1 shows all of the results

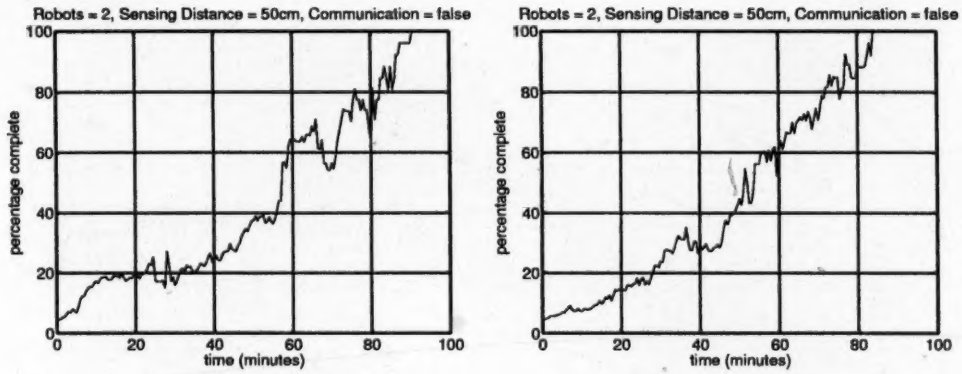

(a) experiment one (b) experiment two

Figure 6.5: These graphs represent the change in percent complete versus time for a two-robot experiment in which the robots have a sensing range of 50 cm and are not able to communicate with each other.

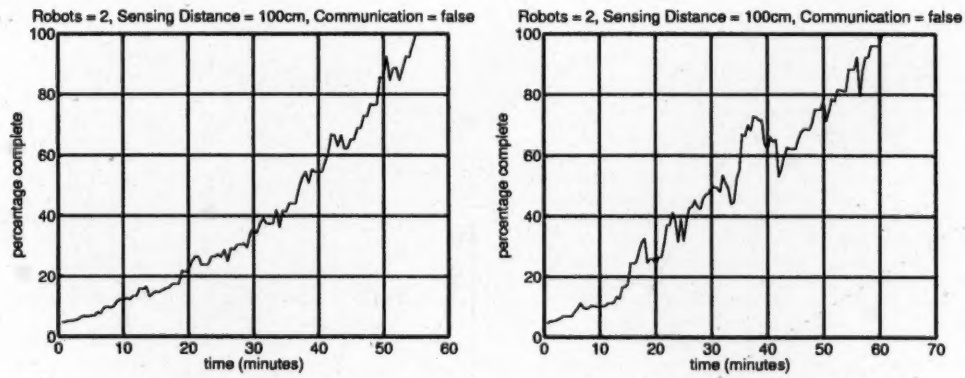

(a) experiment one (b) experiment two

Figure 6.6: These graphs represent the change in percent complete versus time for a two-robot experiment in which the robots have a sensing range of 100 cm and are not able to communicate with each other.

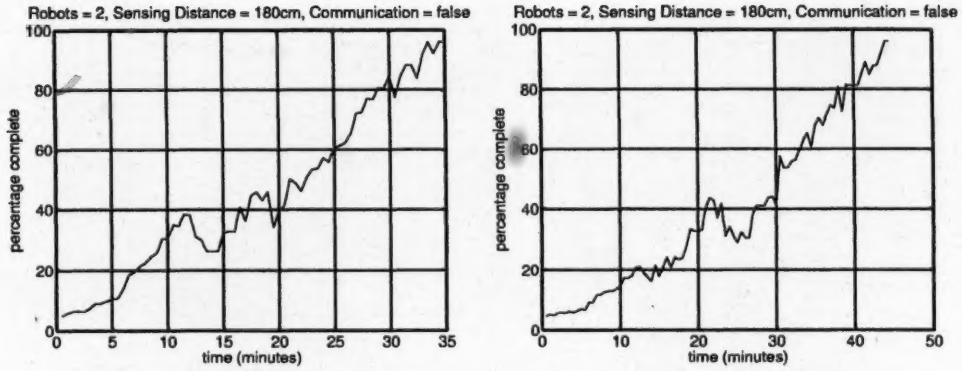
Robots = 2, Sensing Distance = 180cm, Communication = false
percentage complete
time (minutes)
Robots = 2, Sensing Distance = 180cm, Communication = false
percentage complete
time (minutes)

(a) experiment one (b) experiment two

Figure 6.7: These graphs represent the change in percent complete versus time for a two-robot experiment in which the robots have a sensing range of 180 cm and are not able to communicate with each other.

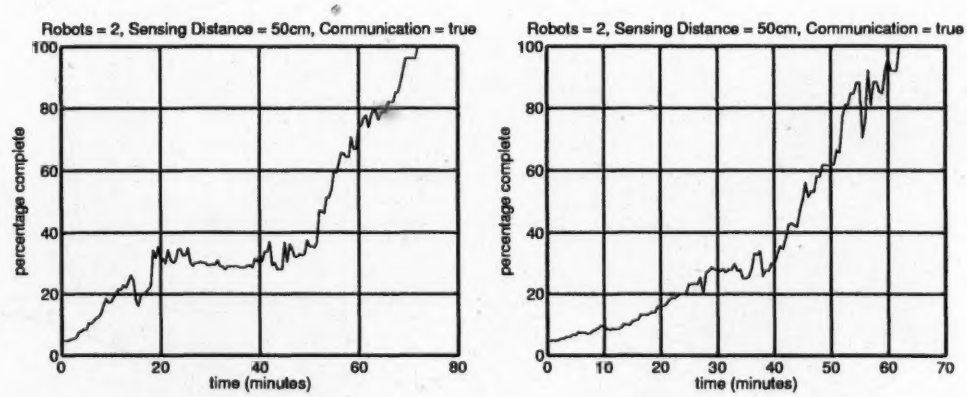
Robots = 2, Sensing Distance = 50cm, Communication = true
percentage complete
time (minutes)
Robots = 2, Sensing Distance = 50cm, Communication = true
percentage complete
time (minutes)

(a) experiment one (b) experiment two

Figure 6.8: These graphs represent the change in percent complete versus time for a two-robot experiment in which the robots have a sensing range of 50 cm and are able to communicate with each other.

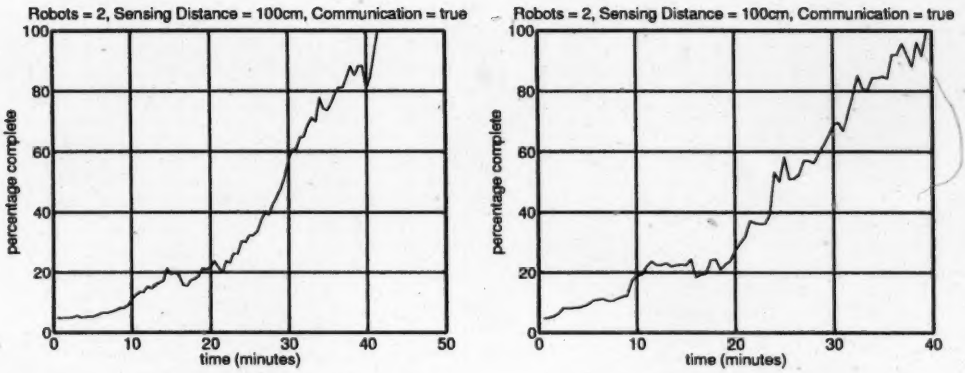

(a) experiment one (b) experiment two

Figure 6.9: These graphs represent the change in percent complete versus time for a two-robot experiment in which the robots have a sensing range of 100 cm and are able to communicate with each other.

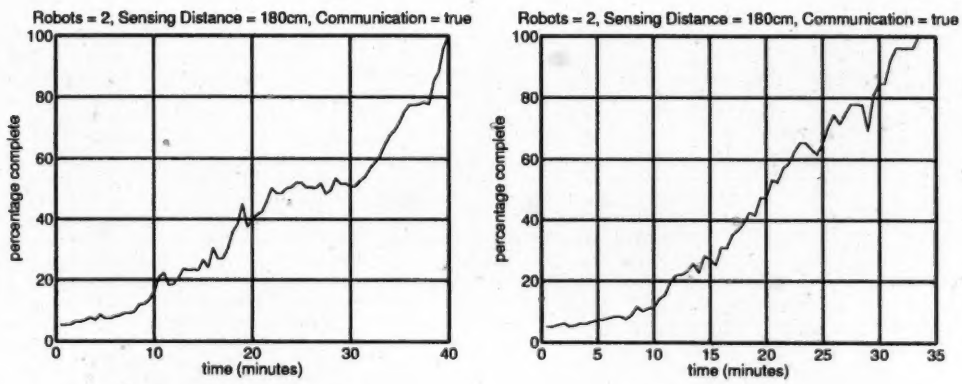

(a) experiment one (b) experiment two

Figure 6.10: These graphs represent the change in percent complete versus time for a two-robot experiment in which the robots have a sensing range of 180 cm and are able to communicate with each other.

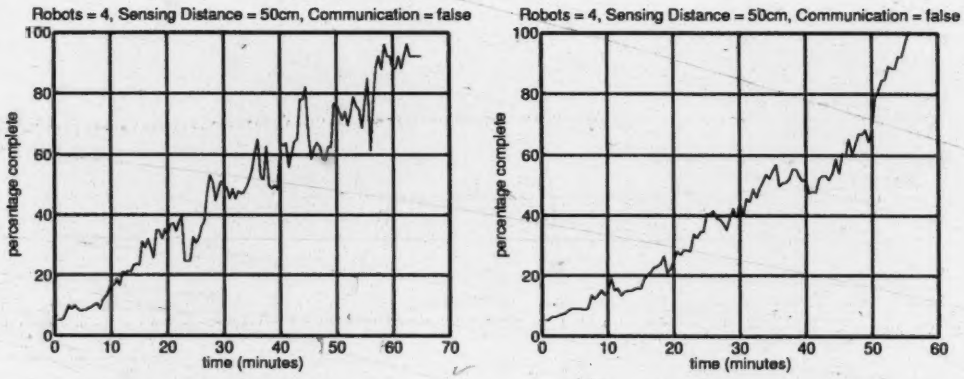
Robots = 4, Sensing Distance = 50cm, Communication = false
percentage complete
time (minutes)
Robots = 4, Sensing Distance = 50cm, Communication = false
percentage complete
time (minutes)

(a) experiment one (b) experiment two

Figure 6.11: These graphs represent the change in percent complete versus time for a four-robot experiment in which the robots have a sensing range of 50 cm and are not able to communicate with each other.

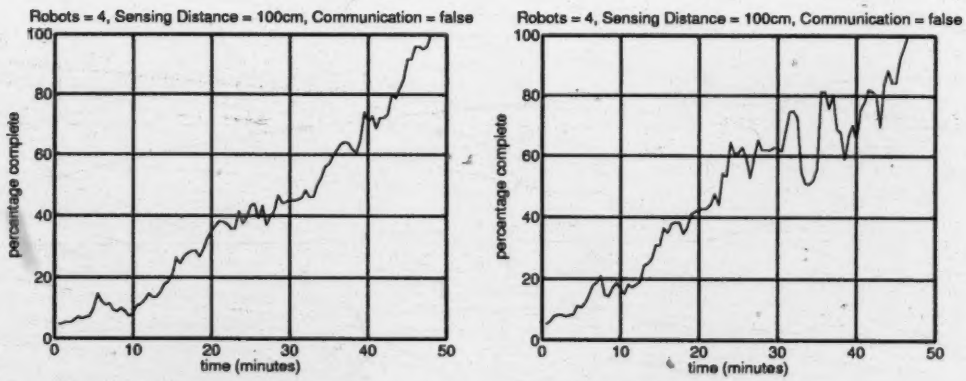
Robots = 4, Sensing Distance = 100cm, Communication = false
percentage complete
time (minutes)
Robots = 4, Sensing Distance = 100cm, Communication = false
percentage complete
time (minutes)

(a) experiment one (b) experiment two

Figure 6.12: These graphs represent the change in percent complete versus time for a four-robot experiment in which the robots have a sensing range of 100 cm and are not able to communicate with each other.

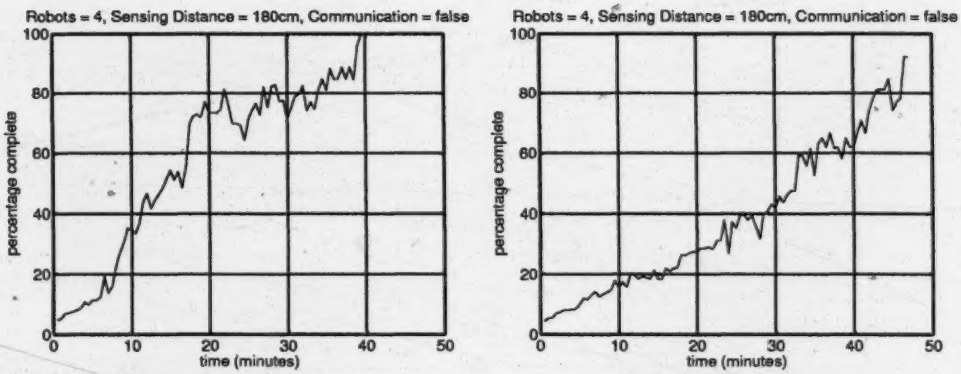
Robots = 4, Sensing Distance = 180cm, Communication = false
percentage complete
time (minutes)
Robots = 4, Sensing Distance = 180cm, Communication = false
percentage complete
time (minutes)

(a) experiment one (b) experiment two

Figure 6.13: These graphs represent the change in percent complete versus time for a four-robot experiment in which the robots have a sensing range of 180 cm and are not able to communicate with each other.

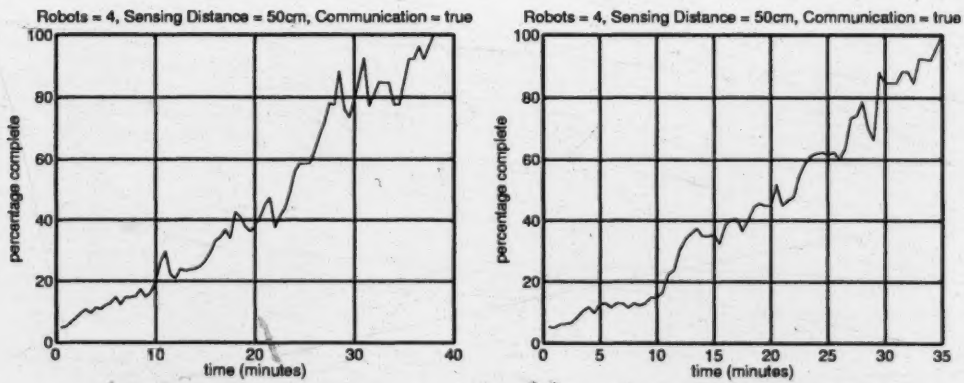
Robots = 4, Sensing Distance = 50cm, Communication = true
percentage complete
time (minutes)
Robots = 4, Sensing Distance = 50cm, Communication = true
percentage complete
time (minutes)

(a) experiment one (b) experiment two

Figure 6.14: These graphs represent the change in percent complete versus time for a four-robot experiment in which the robots have a sensing range of 50 cm and are able to communicate with each other.

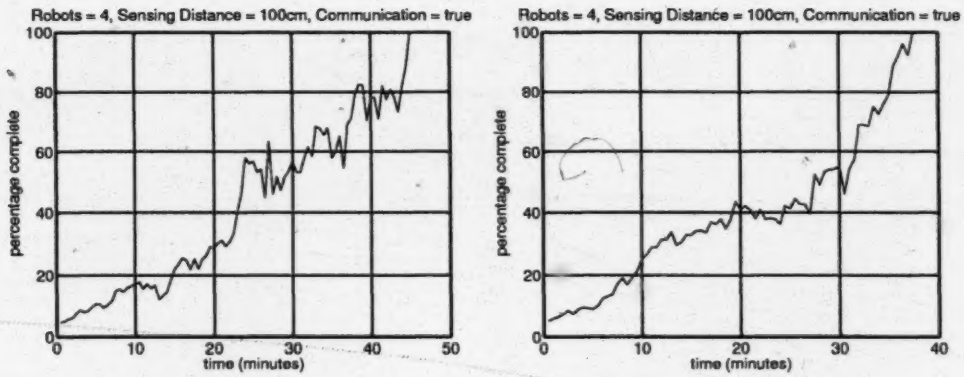
Robots = 4, Sensing Distance = 100cm, Communication = true
percentage complete
time (minutes)
Robots = 4, Sensing Distance = 100cm, Communication = true
percentage complete
time (minutes)

(a) experiment one (b) experiment two

Figure 6.15: These graphs represent the change in percent complete versus time for a four-robot experiment in which the robots have a sensing range of 100 cm and are able to communicate with each other.

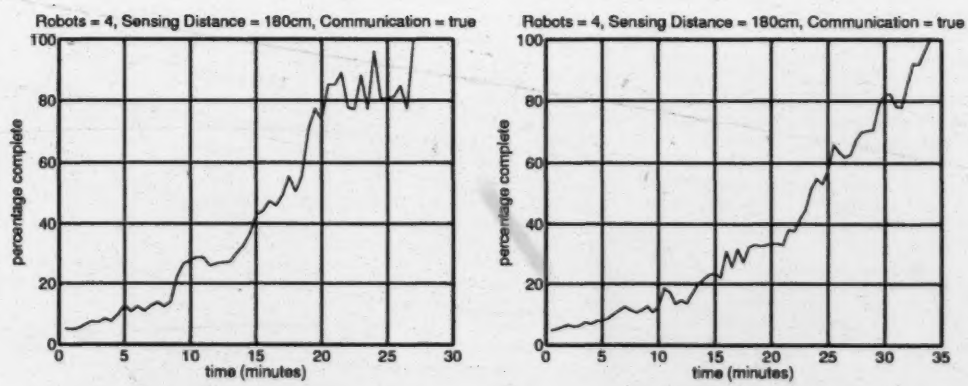
Robots = 4, Sensing Distance = 180cm, Communication = true
percentage complete
time (minutes)
Robots = 4, Sensing Distance = 180cm, Communication = true
percentage complete
time (minutes)

(a) experiment one (b) experiment two

Figure 6.16: These graphs represent the change in percent complete versus time for a four-robot experiment in which the robots have a sensing range of 180 cm and are able to communicate with each other.

6.6 Summary

In this chapter we presented the experimental results for both the single-robot experiments and the multi-robot experiments. The single-robot experiments were presented first to prove the viability of the algorithm shown back in Chapter 4. The multi-robot experiments were then presented demonstrating the improvement on overall system performance when more than one robot is present. A discussion of the results can be found in the next chapter.

Chapter 7

Discussion

7.1 Introduction

In this chapter we will discuss several aspects of the experiments and examine many of the different conditions under which the experiments were performed. We will examine the effects of the initial conditions, the population size, the environment and the sensing and communication range. We will then go on to point out some of the trends that appeared in a majority of the experiments and explain how the stopping criterion was decided upon and carried out. Finally we will compare our experiments with other multi-robot systems.

7.2 Effect of Initial Conditions

The initial conditions of the experiments were outlined in Section 5.5. Figure 5.11 shows a picture of what the system looked like before the experiments. The initial conditions of the system seemed not to affect the overall results, although different initial configurations were not tested. We made this assumption based on the fact that in nearly every experiment, after less than five or ten minutes several of the objects to be sorted had been moved and the world didn't look much at all like the initial state. In each of the experiments shown in Figure 7.1, it is apparent that after less than ten minutes the dispersion of the objects is very different than the initial conditions.

Another thing to note when looking at Figure 7.1 is that small collections of objects are already forming. In nearly all of the experiments the location of these initial larger concentrations proved to be the final location of the finished pile. Figure 7.2 displays the final resulting piles of the initial piles shown in Figure 7.1 and it is easily seen that the initial smaller collections of piles in Figure 7.1 grow to become the larger piles in Figure 7.2.

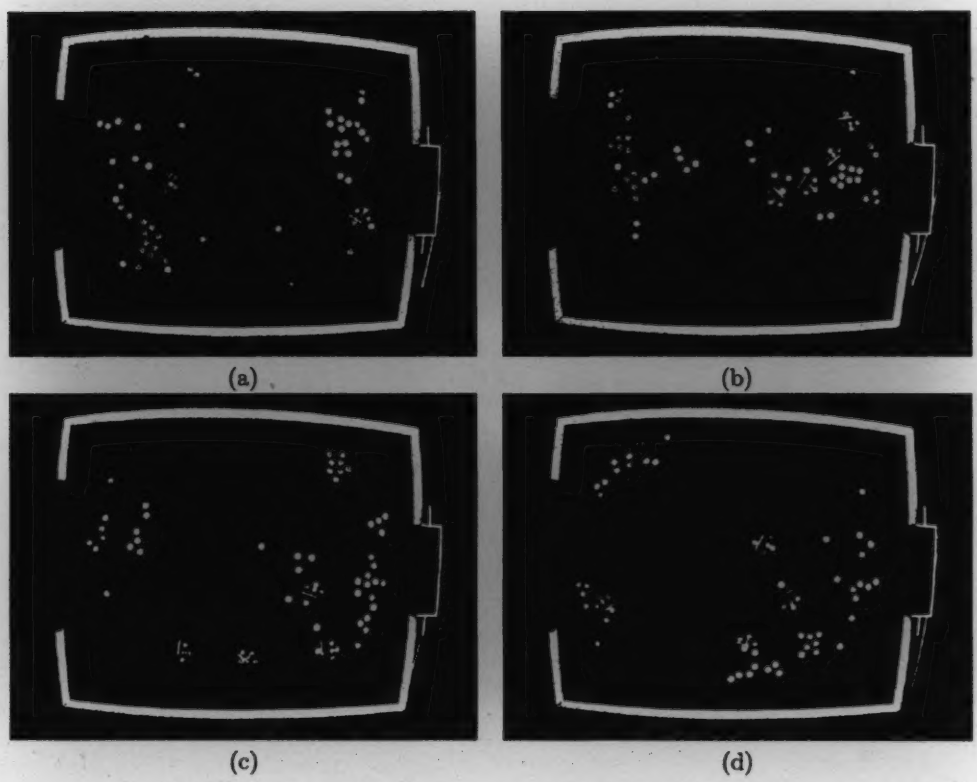

Figure 7.1: Four examples of sorting experiments shortly after the experiment has been started. (a) 2 robots with a sensing range of 180cm and no communication after 8 minutes. (b) 4 robots with a sensing range of 100cm and no communication after 8 minutes. (c) 4 robots with a sensing range of 100cm and with communication after 9 minutes. (d) 4 robots with a sensing range of 50cm and with communication after 10 minutes. Note: There are two types of objects (yellow and orange) as in Figure 5.11, except that yellow objects have been painted white with an image editor to improve the readability of the images.

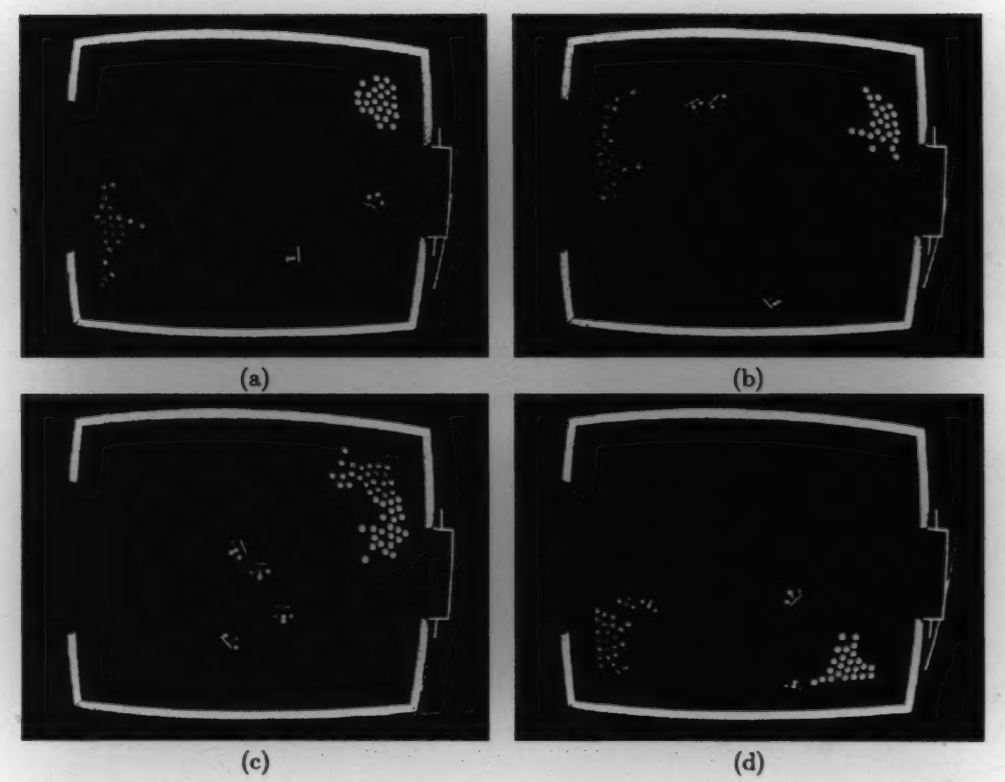

Figure 7.2: Four examples of sorting experiments after the experiment have completed. (a) 2 robots with a sensing range of 180cm and no communication. (b) 4 robots with a sensing range of 100cm and no communication. (c) 4 robots with a sensing range of 100cm and with communication. (d) 4 robots with a sensing range of 50cm and with communication. Note: There are two types of objects (yellow and orange) as in Figure 5.11, except that yellow objects have been painted white with an image editor to improve the readability of the images.

7.3 Effect of Population Size

Examining Table 6.1 and Figure 7.3, shows that adding more robots to the environment increases the overall efficiency of the system. However, adding additional robots to the system did not linearly increase the efficiency. I.e., if we doubled the number of robots that did not translate into half the total time needed to complete the task. The reason that the task does not get completed in half the time is because the environment now contains more robots and therefore each robot must spend more time avoiding other robots. Goldberg and Matarić [25] based experiments solely on the idea of using interference between robots as a key diagnostic parameter for measuring the efficiency of a multi-robot system. While we did not keep track of the amount of time that each robot was avoiding other robots, our experiments have proven that while increasing population size does increase the overall efficiency of the system, this improvement is not linear with respect to the number of robots added. Another factor regarding the effect population size has on the environment is the overall environment size. The next section deals with that specific effect.

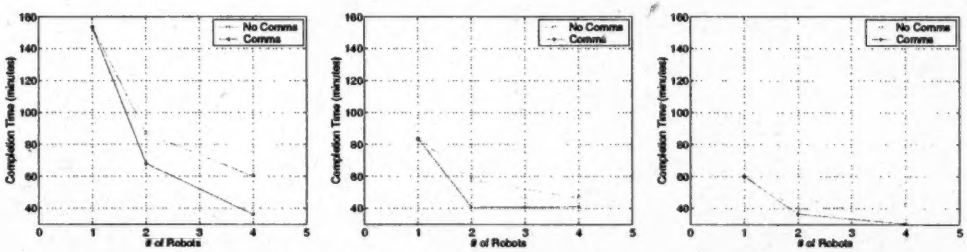

(a) 50 cm sensing range (b) 100 cm sensing range (c) 180 cm sensing range

Figure 7.3: Completion times for each of the three robot numbers with 180cm sensing range when there is either communication (bottom curve) or no communication (top curve) between the robots.

7.4 Effect of Environment

We believe that both the completion rate and completion efficiency are dependent on the environment. The robots' environment includes the area used for sorting, the other robots, and the objects being sorted. The sizes and shapes of all of these objects affected our results. Tests were not conducted with a larger sorting area or smaller robots, but it is likely that, if the facilities were available to run tests varying these environmental elements we would find that an increased area and smaller robots would produce less interference between robots and thus perhaps even more robots could be added to the system.

Each robot was 0.18 m in diameter and thus took up 0.025 m^2 of the field. The sorting area was 1.8 m×2.8 m in size and thus took up 5.04m^2. Therefore each robot took up

0.5% of the entire area. By increasing the environment by 0.75 m on each side, essentially doubling the area of the environment, each robot would take up 0.25% of the entire area. This would give it twice as much space to operate and result in less interference situations with another robot. The same effect could be seen by keeping the same sorting area and reducing the robot diameter to 0.12 m.

Unfortunately, we were not able to conduct any experiments like these. However, the point we are making is that the environment used in our experiments indeed had an effect on our results.

7.5 Effect of Sensing and Communication Range

7.5.1 Sensing Range

The sensing ranges chosen for our experiments were 50 cm, 100 cm, and 180 cm. These ranges were chosen for two specific reasons. In the initial testing stages sensing ranges were going to start be set at 25 cm and increment by 25 cm up to 175 cm. After initial tests, it was shown that a 25 cm range didn't offer enough information to the robot to allow it to make any informed decisions, thus the baseline searching range was changed to 50 cm. After more tests were observed, it was seen that insignificant amounts of change happened between 50 cm and 75 cm and between 75 cm and 100 cm, thus the 75 cm testing range was not used. Similarly the ranges after 100 cm were fairly close in their results and eventually 180 cm was chosen because it was equal to the distance on one side of the sorting area.

As mentioned in Section 6.2, results (Table 6.1) showed that while there was a significant improvement when the sensing range increased from 50 cm to 100 cm, the improvement became much less pronounced when the sensing range was changed from 100 cm to 180 cm. We came up with three possible reasons why this may have occurred. First, this may indicate that the algorithm used by the robot could not take advantage of the additional information beyond a certain sensing threshold. Second, this may also indicate that the difference could be related to Section 7.4, i.e., the environment size. Perhaps a larger environment would have produced better results for the larger sensing ranges. Finally, it's possible that the discrepancy could just be due to the local nature of the sorting task.

7.5.2 Communication Range

For all of the experiments that involved communication, the area that defined the communication range was identical to the area that defined the sensing range. The communication range was chosen like this for one main reason. In previous experiments involving communication, for example [6], LEDs were placed on top of robots signifying state information. This information was easily detected by other robots. Our experiments operated similarly. As in Figure 4.2 if Robot 1 could see Robot 2 then the information Robot 2 had perceived was

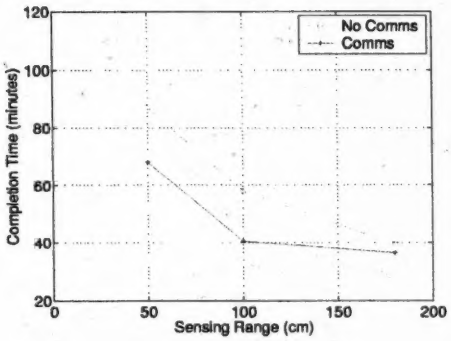

Figure 7.4: Completion times for each of the three sensing ranges with two robots when there is either communication (bottom curve) or no communication (top curve) between the robots

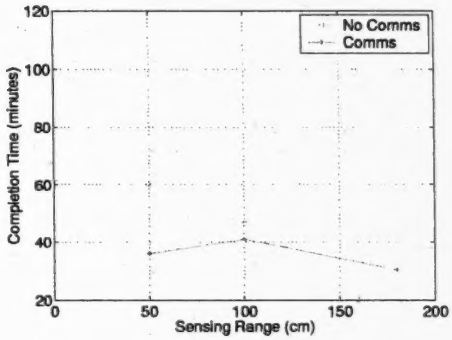

Figure 7.5: Completion times for each of the three sensing ranges with four robots when there is either communication (bottom curve) or no communication (top curve) between the robots

automatically transmitted to Robot 1. This is similar to Robot 2 having the information available on top of the robot and Robot 1 being able to "see" it. It is also important to note that the transfer of information is one way and not two way. This too mimics previous experiments in the way that the robot must be able to "see" just exactly who it is receiving information from.

Another point we found interesting regarding the communication range was that the information we communicated was only useful if the the algorithms running on the robot could handle the information. For example, many other pieces of information could have been communicated (e.g. robot velocity) but if this information isn't useful then there is no reason to communicate it.

7.6 Experiment Specific Tendencies

This section points out two specific events that happened in many of the experiments but might not have been noticed. The time histories of two of the experiments, both involving four robots and a 100 cm sensing range, but one with communication and the other without, are plotted in Figures 7.6 and 7.7. The figures show how the percent completion rate, defined by Equation 4.14, changed with time.

The first graph, Figure 7.6, shows a very orderly convergence toward completion with a sharp increase in overall performance near the end of the experiment. This sharp increase is common in several of the graphs (shown fully in Section 6.5) and is the result of robots having fewer piles to drop objects on to and thus objects are being dropped more and more often on the pile with the largest concentration of similar objects.

The second graph, Figure 7.7, also shows an orderly convergence toward completion, but

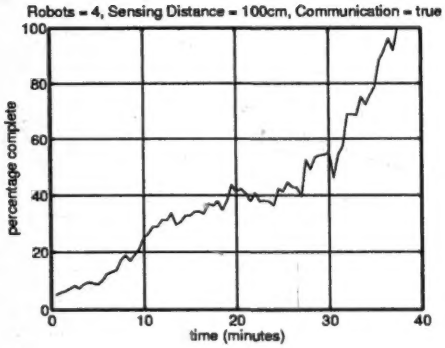

Figure 7.6: This graph represents the change in percent complete versus time for a four-robot experiment in which the robots have a sensing range of 100 cm and are able to communicate with each other. Notice the sharp increase in the slope of the curve at the end of the graph. Several experiments showed results like this because there were fewer piles to drop objects onto.

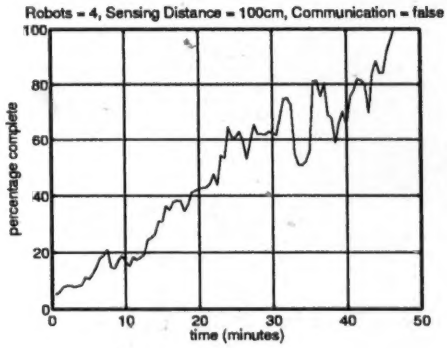

Figure 7.7: This graph represents the change in percent complete versus time for a four-robot experiment in which the robots have a sensing range of 100 cm and are not able to communicate with each other. Notice the dip in performance around the 33 minute mark. This indicates a time when a robot made an incorrect decision and accidentally separated a pile of objects into more than one pile.

encounters a sharp decline in progress just past the 30-minute mark. This sharp decline can be explained by a robot attempting to pick up an object that perhaps it shouldn't, and resulted in a pile of objects being split into two or more piles. However, the system quickly converged back onto its path toward completion once the multiple piles were again joined.

7.7 The Stopping Criterion

The completion percentage evaluation (Equation 4.14) was created as a simple way to evaluate the system. Equation 4.14 always has a value between 0 and 100 and is easy to calculate. The stopping criterion was computed by an outside agent that analyzed the vision data of the entire world. The simple duties of this agent included notifying the operator that the 90% criterion had been met, and keeping track of the overall progress of the experiment. During the course of the experiments we would wait and see if the system would achieve 100% completion but if 100% completion was not achieved in a short amount of time the experiment was stopped and still deemed successful. It should be noted that the individual agents did not have any stopping criterion built into their system. Thus, if there was no "global" agent running the robots would have continued to sort *ad infinitum*. However, if the objects were already fully sorted, the probability the robot would find a suitable object to pick up would be very low and the only possible place to drop the object would be back on the pile it had been picked up from.

7.8 Similarities to Other Multi-Robot Systems

Similarities exist between our results and other multi-robot systems. For instance, in experiments by Balch and Arkin [6], Dahl *et al.* [15], Easton and Martinoli [18] and Matarić [43] it was experimentally shown that communication can improve the overall task performance in a multi-robot system. We also provided a sorting algorithm for segregation sorting and developed this algorithm based on the previous algorithms by Melhuish *et al.* [46] and Deneubourg *et al.* [16]. Finally, multi-robot systems using similar communication and perception strategies are being developed for military and space applications.

7.9 Summary

In this chapter we discussed several aspects of our experiments and examined many of the different conditions under which the experiments were performed. We examined the effects of the initial conditions, the population size, the environment and the sensing and communication range. We pointed out some of the tendencies that occurred in a majority of the experiments and explained how the stopping criterion was decided upon and carried out. Finally we noted the similarities of our experiments with other multi-robot systems. In the next chapter we will present our conclusions and discuss possible future work.

Chapter 8

Conclusions

8.1 What has been done?

We can make several important observations about the experimental results we have obtained and described in the previous chapters. First we have created a successful segregation sorting algorithm. Second, we have provided more experimental evidence that communication in multi-robot systems improves task performance. Finally, we have developed a relationship between communication and perception in multi-robot systems.

8.1.1 Successful Segregation Sorting Algorithm

We believe our algorithm is the first that has successfully demonstrated segregation sorting *experimentally* given a completion criterion of greater than 90% for Equation 4.14. In fact, of the 30 experiments conducted, our system was able to achieve 100% completion in terms of Equation 4.14 in 25 trials, with an average completion rate of 99%! 100% completion did not occur in an experiment only when, for example, a Type 1 object was stranded in the middle of a large pile of Type 2 objects, rendering it unreachable by the robots (see Figure 8.1). We concluded that a more sophisticated control strategy needs to be developed in order to overcome this problem.

8.1.2 Communication Improves Task Performance

The use of communication between robots in the sorting task generally improves the task performance, a result that agrees with the literature on communication in other collective robotics tasks. This is clearly demonstrated in plots of the data in Table 6.1 and in Figures 7.4 and 7.5, for the case of two-robot and four-robot experiments, respectively. In all cases, communication improved the performance of the tasks. This improvement was particularly pronounced when the sensing range was small. It is also interesting to note that communication had a reduced positive effect on task performance as the sensing range increased. This is quite understandable as a robot with a large sensing range may gain

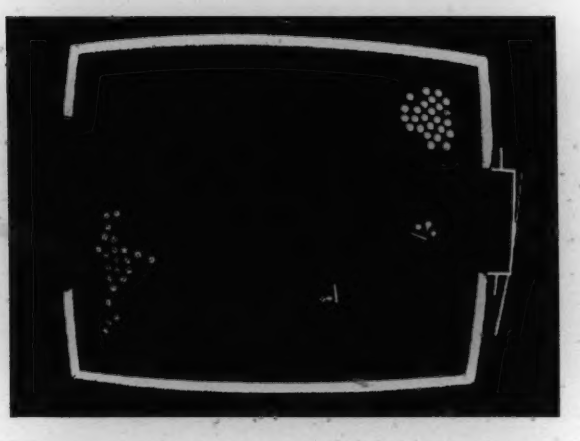

Figure 8.1: As can been seen in this image, one of the objects of type 1 has been stranded in the middle of the type 2 objects. The completion criteria was still satisfied, however, the experiment didn't finish with a 100% completion rate. This experiment involved 2 robots with a sensing range of 180 cm and no communication between the robots.

relatively little additional information from communication if it can already perceive much of what it needs to make a sound decision.

8.1.3 Communication is Related to Perception

Perhaps the most interesting observation is that we have produced empirical evidence to support the argument that there is an equivalence relationship between sensing and communication; i.e., reduced sensing can be compensated for by increased communication. For example, in the 2-robot experiments shown in Figure 7.4, using completion time as the metric, the system with a 180 cm sensing range and no communication is equivalent to the system with a 100 cm sensing and communication range; both achieve a 40-minute average completion time. This result is both intuitive and useful for making decisions regarding design trade-offs in controller configuration.

8.2 Importance

The research we've performed has several significant aspects. We have proven that simple robots with basic behaviours are able to accomplish a task as a group and are able to interact with each other to accomplish this task. This is important because collective robotics is a relatively new area of research and any successful implementation of a multi-robot system only enhances the credibility of this area of research. We have also proven that communication of data between robots improves the overall task performance. This is important because it proves that, assuming systems are able to handle and interpret the information communicated, multi-robot systems can communicate and operate more efficiently when

they are sharing information with each other. Finally, we have proven that there is a relationship between communication and perception. We feel this is important because if we can realize this connection, then we can further develop algorithms that are able to handle the various types of information that may be communicated.

Finally, we have again proven that a multi-robot system is able to adapt to an ever changing environment, even when the robot is changing the environment itself.

8.3 Future Research

Many different facets of research can evolve from this research. Possible future work may include the development of theoretical support for our empirical observations, and the investigation of multi-class collective sorting. Other future work might involve different aspects of communication, or adding extra elements to assist communication (e.g., memory). We also believe that different methods of communication will be researched attempting to further improve the efficiency of multi-robot systems. Results from this work should also help in the fields of multi-robot reconnaissance and map-building missions. We believe that additional research will also be done asking the question, can increased communication or perception harm the system's efficiency?

Bibliography

[1] Carl Anderson and Nigel R. Franks. Teams in animal societies. *Behavioral Ecology*, 12(5):534–540, 2001.

[2] Tamio Arai, Enrico Pagello, and Lynne E. Parker. Editorial: Advances in multi-robot systems. *IEEE Transactions on Robotics and Automation*, 18(5):655–661, October 2002.

[3] Ronald C. Arkin and Tucker Balch. Cooperative multiagent robotic systems. In David Kortenkamp, R.P. Bonasso, and R. Murphy, editors, *Artificial Intelligence and Mobile Robots*, Cambridge, MA, 1998. MIT/AAAI Press.

[4] Ronald C. Arkin and Johnathan Diaz. Line-of-sight constrained exploration for reactive multiagent robotic teams.

[5] Tucker Balch. *Robot Teams: From Diversity to Polymorphism*. A. K. Peters, Ltd., 2002. (ISBN: 1-56881-155-1).

[6] Tucker Balch and Ronald C. Arkin. Communication in reactive multiagent robotic systems. *Autonomous Robots*, 1(1):27–52, 1994.

[7] Tucker Balch and Ronlad C. Arkin. Behavior-based formation control for multi-robot teams. *IEEE Transactions on Robotics and Automation*, 1999.

[8] Tucker R. Balch, Gary Boone, Thomas R. Collins, Harold Forbes, Doug MacKenzie, and Juan Carlos Santamaria. Io, ganymede, and callisto: A multiagent robot trash-collecting team. *AI Magazine*, 16(2):39–51, 1995.

[9] D. Barnes and J. Gray. Behaviour synthesis for cooperant mobile robot control. In *International Conference on Control*, pages 1135–1140, 1991.

[10] R. Beckers, O.E. Holland, and J.L. Deneubourg. From local actions to global tasks: Stigmergy and collective robotics. In R. Brooks and P. Maes, editors, *Artifical Life IV*. MIT Press, 1994.

[11] Valentino Braitenberg. *Vehicles: Experiments in Synthetic Psychology*. MIT Press, Cambridge MA, 1984.

[12] Rodney A. Brooks. A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, RA-2(1):14–23, March 1986.

[13] Rodney A. Brooks. From earwigs to humans. In *Proceedings IIAS The Third Brain and Mind International Symposium Concept Formation, Thinking and Their Development*, pages 59–66, Kyoto, Japan, May 1996.

[14] Y. Uny Cao, Alex S. Fukunaga, and Andrew B. Kahng. Cooperative mobile robotics: Antecedents and directions. *Autonomous Robots*, 4(1):7–23, March 1997.

[15] Torbjorn S. Dahl, Maja J. Matarić, and Gaurav S. Sukhatme. Adaptive spatio-temporal organization in groups of robots. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'02)*, pages 1044–1049, Lausanne, Switzerland, September 2002.

[16] J.L. Deneubourg, S. Goss, N. Franks, A. Sedova-Franks, C. Detrain, and L. Chretien. The dynamics of collective sorting: Robot-like ants and ant-like robots. In J.A. Meyer and S.W. Wilson, editors, *Simulation of Adaptive Behavior*, Animals to Animats, 1991.

[17] Gregory Dudek, Michael R. M. Jenkin, Evangelos Milios, and David Wilkes. A taxonomy for multi-agent robotics. *Autonomous Robots*, 3:375–397, 1996.

[18] Kjerstin I. Easton and Alcherio Martinoli. Efficiency and optimization of explicit and implicit communication schemes in collaborative robotics experiments. In *IEEE Conference on Intelligent Robots and Systems IROS-02*, pages 2795–2800, Lausanne, Switzerland, September 2002. IEEE Press.

[19] Donald E. Franklin, Andrew B. Kahng, and M. Anthony Lewis. Distributed sensing and probing with multiple search agents: toward system-level landmine detection solutions. In *Detective Technologies for Mines and Minelike Targets, Proceedings of SPIE*, volume 2496, pages 698–709, 1995.

[20] Nigel R. Franks and Jean-Louis Deneubourg. Self organizing nest construction in ants: individual worker behaviour and the nest's dynamics. *Animal Behaviour*, 54:779–796, 1997.

[21] N.R. Franks, A. Wilby, B.W. Silverman, and C. Tofts. Self organizing nest construction in ants: sophisticated building by blind bulldozing. *Animal Behaviour*, 44:357–375, 1992.

[22] Jakob Fredslund and Maja J Matarić. Robot formations using only local sensing and control. In *IEEE International Symposium on computationl Intelligence for Robotics and Automation (CIRA-2001)*, Banff, Canada, July 2001.

[23] Douglas W. Gage. Command control for many-robot systems. *Unmanned Systems*, 10(4):28–34, Fall 1992.

[24] J. Gautrais, G. Theraulaz, J.L. Deneubourg, and C. Anderson. Emergent polyethism as a consequence of increased colony size in insect societies. *Journal of Theoretical Biology*, 215:363–373, 2002.

[25] Dani Goldberg and Maja J Matarić. Interference as a tool for designing and evaluating multi-robot controllers. In *AAAI-97*, pages 637–642, Providence, Rhode Island, July 1997.

[26] Robert Grabowski, Luis E. Navarro-Serment, Christiaan J. J. Paredis, and Pradeep K. Khosla. Heterogeneous teams of modular robots for mapping and exploration. *Autonomous Robots*, 8(3):293–308, 2000.

[27] Alan R. Graves and Chris Czarnecki. Design patterns for behavior-based robotics. *IEEE Transactions on Systems Man and Cybernetics Part A: Systems and Humans*, 30(1):36–41, January 2000.

[28] Point Grey. Point grey official site, 2004. http://www.ptgrey.com/.

[29] Thomas C. Hales. The honeycomb conjecture. *Discrete and Computational Geometry*, 1999.

[30] Owen Holland and Chris Melhuish. Stigmergy, self-organization, and sorting in collective robotics. *Artificial Life*, 5:173–202, 1999.

[31] Bert Hölldobler and Edward O. Wilson. *The Ants*. HUP, Cambridge, MA, 1990.

[32] T. L. Huntsberger, M. Matarić, and P. Pirjanian. Action selection within the context of a robotic colony. In *SPIE Symposium on Sensor Fusion and Decentralized Control in Robotic Systems II*, volume 3839, pages 84–90, Boston, MA, September 1999.

[33] T. L. Huntsberger, G. Rodriguez, and P. S. Schenker. Robotics challenges for robotic and human mars exploration. In *ROBOTICS2000*, pages 340–346, Albuquerque, NM, March 2000.

[34] T. L. Huntsberger and J. Rose. Behavior-based control for autonomous mobile robots. In *ROBOTICS2000*, pages 299–305, Albuquerque, NM, March 2000.

[35] C. Ronald Kube and Eric Bonabeau. Cooperative transport by ants and robots. *Robotics and Autonomous Systems*, 30(1/2):85–101, 2000. ISSN: 0921-8890.

[36] C. Ronald Kube and Hong Zhang. Collective robotics: From social insects to robots. *Adaptive Behavior*, 2(2):189–218, 1994.

[37] Kristina Lerman and Aram Galstyan. Mathematical model of foraging in a group of robots: Effect of interference. *Autonomous Robots*, 13(2):127–141, 2002.

[38] Pattie Maes and Rodney A. Brooks. Learning to coordinate behaviors. In *National Conference on Artificial Intelligence*, pages 796–802, 1990.

[39] M. Maris and R. Boekhorst. Exploiting physical constraints: heap formation through behavioral error in a group of robots. In *Proc. IROS'96, IEEE/RSJ International Conference on Intelligent Robots and Systems (in press)*, page 20, 1996.

[40] A. Martinoli, A. J. Ijspeert, and F. Mondada. Understanding collective aggregation mechanisms: From probabilistic modelling to experiments with real robots. *Special Issue on Distributed Autonomous Robotic Systems, Robotics and Autonomous Systems*, 29:51–63, 1999.

[41] A. Martinoli and F. Mondada. Collective and cooperative group behaviours: Biologically inspired experiments in robotics. In O. Khatib and J.K. Salisbury, editors, *Proc. of the Fourth Int. Symp. on Experimental Robotics ISER-95*, Stanford US, June-July 1995. Lecture Notes in Control and Information Sciences (1997), Vol. 223, pp. 3-10. Springer Verlag.

[42] Maja J. Matarić. *Interaction and Intelligent Behavior*. PhD thesis, MIT, EECS, May 1994.

[43] Maja J. Matarić. Using communication to reduce locality in distributed multi-agent learning. *Journal of Experimental and Theoretical Artificial Intelligence*, 10(3):357–369, Jul-Sep 1998. special issue on Learning in DAI Systems, Gerhard Weiss, ed.

[44] Maja J. Matarić, Martin Nilsson, and Kristian T. Simsarian. Cooperative multi-robot box-pushing. In *IEEE/RSJ IROS*, pages 556–561, 1995.

[45] Matthew McNaughton and Hong Zhang. Color vision for robocup with fast lookup tables. In *IEEE International Conference on Robotics, Intelligent Systems and Signal Processing*, October 8-13 2003.

[46] Chris Melhuish, Owen Holland, and Steve Hoddell. Collective sorting and segregation in robots with minimal sensing. In *5th Int. Conf. on the Simulation of Adaptive Behaviour*, 1998.

[47] Merriam-Webster. *Webster's 7th Collegiate Dictionary*. Merriam-Webster Inc., 1963.

[48] Mark W. Moffet. Cooperative food transport by an asiatic ant. *National Geographic Research*, 4(3):386–394, 1988.

[49] Hans Mori, Yutaka Saito, and Yow Pong Tho. Cooperative group predation in a sit and wait cheyletid mite. *Experimental and Applied Acarology*, 23:643–651, 1999. Kluwer Academic Publishers.

[50] Robin R. Murphy. *Introduction to AI Robotics*. A Bradford Book, 2000.

[51] Chris A. C. Parker, Hong Zhang, and C. Ronald Kube. Blind bulldozing: Multiple robot nest construction. In *IEEE Conference on Intelligent Robots and Systems IROS-03*, 2003.

[52] Christopher Parker and Hong Zhang. Robot collective construction by blind bulldozing. In *IEEE Conference on Systems Cybernetics and Man*, 2002.

[53] Lynne E. Parker. An experiment in mobile robotic cooperation. In *Proceedings of the ASCE Specialty Conference on Robotics for Challenging Environments*, pages 131–139, February 1994.

[54] Lynne E. Parker. *Distributed Autonomous Robotic Systems 4*, chapter Current State of the Art in Distributed Robot Systems, pages 3–12. Springer, 2000.

[55] S. Premvuti and S. Yuta. Consideration on the cooperation of multiple autonomous mobile robots. In *IEEE International Workshop of Intelligent Robots and Systems*, pages 59–63, Tsuchiura, Japan, 1990.

[56] Radiometrix. Rpc-418-40 official site, 2004. http://www.radiometrix.co.uk/.

[57] Craig W. Reynolds. Flocks, herds, and schools: A distributed behavioral model. *Computer Graphics*, 21(4):25–34, July 1987.

[58] RoboCup. Robocup official site, 2004. http://www.robocup.org/.

[59] RoboMinds. Robominds official site, 2004. http://www.robominds.com/.

[60] Dunbar W. S. and Klein B. Mining, mineral processing, and mini-machines. *CIM Bulletin*, 95(1057), January 2002.

[61] W. Tang and H. Zhang. Decentralized control of robot formation marching. In *Sixth International Symposium on Robotics and Manufacturing*, Montepellier, France, May 27-30 1996.

[62] Ashitey Trebi-Ollennu, Hari Das Nayar, Hrand Aghazarian, Anthony Ganino, Paolo Pirijanian, Brett Kennedy, Terry Huntsberger, and Paul Schenker. Mars rover pair cooperatively transporting a long payload. In *IEEE International Conf. on Robotics and Automation (ICRA)*, pages 3136–3141, Washington, DC, May 2002.

[63] W.G. Walter. An imitation of life. *Scientific American*, 182(5):42–45, May 1950.

[64] W.G. Walter. A machine that learns. *Scientific American*, 185(2):60–63, August 1950.

[65] Jens Wawerla, Gaurav S. Sukhatme, and Maja J. Matarić. Collective construction with multiple robots. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2696–2701, Lausanne, Switzerland, 2002.